

User Guide
for
OmniVista 2500 NMS Enterprise
Version 4.2.1.R01
Maintenance Release 2



May 2017
Revision A
READ THIS DOCUMENT
ALE USA Inc.
26801 West Agoura Road
Calabasas, CA 91301
+1 (818) 880-3500

Table of Contents

1.0 Getting Started with OmniVista 2500 NMS	1-1
OmniVista 2500 NMS User Interface	1-1
Applications	1-2
Customizing the Dashboard	1-3
Configuration/Display Buttons	1-4
Working with Tables	1-6
Setting Up OmniVista 2500 NMS for Network Management	1-7
Discovering Devices	1-7
Editing Discovered Devices	1-7
Why Edit a Discovered Device?	1-8
Configuring Traps	1-9
Saving Changes	1-9
PolicyView QoS	1-10
2.0 Dashboard	2-1
Overview	2-1
Viewing Additional Information	2-2
Configuring the Widget Display	2-3
Linking to the Application	2-3
Customizing the Dashboard	2-4
Adding Widgets	2-4
Removing Widgets	2-5
Changing the Dashboard Layout	2-5
Application Widgets	2-6
3.0 Analytics	3-1
Using Analytics	3-1
Configuring Analytics	3-3
Reports	3-3
Report Options	3-4
Top N Applications	3-10
Top N Applications - Advanced	3-17
Top N Clients	3-24
Top N Switches	3-29
Top N Ports Utilization	3-31
Network Availability	3-36
Alarms	3-37
Active Calls	3-39
Ended Calls	3-41
Performance Monitoring	3-43
Using the Performance Monitoring Feature	3-44
Viewing Statistics Profile Data	3-48
Editing a Statistics Profile	3-53
Deleting a Statistics Profile	3-53
Scheduling a Profile	3-53
Starting/Stopping a Profile	3-54
Statistics Profile List	3-54

Table of Contents (continued)

Profiles	3-54
Creating a Profile	3-55
Editing a Profile	3-56
Deleting a Profile	3-56
Viewing a Profile	3-56
Summary View	3-56
Applications Management	3-57
Creating Application Mapping	3-57
Editing Application Mapping	3-58
Deleting Application Mapping	3-58
Anomalies.....	3-58
Settings	3-58
Analytics Configuration	3-58
Application Visibility Statistics Configuration.....	3-59
4.0 App Launch	4-1
Edit Mode	4-1
5.0 Application Visibility.....	5-1
Application Visibility Application	5-1
Application Visibility Configuration	5-2
Devices Management.....	5-3
Device Information.....	5-3
Signature Files	5-4
Viewing Signature Files	5-4
Importing a Signature File.....	5-4
Upgrading a Signature File	5-4
Deleting a Signature File.....	5-5
Signature Profiles.....	5-5
Viewing Signature Profiles.....	5-5
Creating a Signature Profile.....	5-5
Editing a Signature Profile	5-9
Cloning a Signature Profile	5-9
Applying a Signature Profile	5-9
Removing a Signature Profile	5-10
Deleting a Signature Profile	5-10
Summary View	5-10
Settings	5-11
Audit Configuration	5-11
Update Configuration.....	5-11
6.0 Audit.....	6-1
Log Files by Type	6-1
Viewing Log Files	6-2
Searching Through Log Files.....	6-3
Filtering Log File Entries	6-3
Downloading Log Files	6-4
All Current Logs.....	6-4
Settings	6-4

Table of Contents (continued)

7.0 Authentication Servers	7-1
LDAP Servers.....	7-1
RADIUS Servers	7-2
ACE Servers.....	7-2
TACACS+ Servers	7-2
LDAP Server Management.....	7-2
Adding an LDAP Server.....	7-3
Modifying an LDAP Server	7-4
Deleting an LDAP Server.....	7-4
Configuring an LDAP Server	7-4
LDAP Server Management Table.....	7-5
RADIUS Server Management	7-6
Adding a RADIUS Server.....	7-6
Modifying a RADIUS Server	7-7
Deleting a RADIUS Server.....	7-7
Configuring a RADIUS Server	7-7
RADIUS Server Management Table.....	7-7
ACE Server Management.....	7-8
Adding an ACE Server.....	7-8
Deleting an ACE Server.....	7-8
Configuring an ACE Server	7-8
TACACS+ Server Management	7-9
Adding a TACACS+ Server	7-9
Modifying a TACACS+ Server	7-10
Deleting a TACACS+ Server	7-10
TACACS+ Server Management Table	7-10
8.0 Captive Portal	8-1
Configuration.....	8-1
Creating a Captive Portal Configuration	8-1
Editing a Captive Portal Configuration.....	8-2
Assigning a Captive Portal Configuration	8-2
Deleting a Captive Portal Configuration.....	8-2
Profile	8-2
Creating a Captive Portal Profile	8-3
Editing a Captive Portal Profile.....	8-3
Assigning a Captive Portal Profile	8-3
Deleting a Captive Portal Profile.....	8-3
Profile Domain Policy List.....	8-3
Creating a Profile Domain Policy	8-4
Editing a Profile Domain Policy.....	8-4
Assigning a Profile Domain Policy	8-4
Deleting a Profile Domain Policy	8-4
Domain Policy List.....	8-4
Creating a Captive Portal Domain Policy List	8-4
Editing a Captive Portal Domain Policy List.....	8-5
Assigning a Captive Portal Domain Policy List	8-5
Deleting a Captive Portal Domain Policy List.....	8-5

Table of Contents (continued)

Captive Portal Customization	8-5
Creating a Captive Portal Customization	8-5
Editing a Captive Portal Custom File	8-6
Applying a Captive Portal File	8-6
Deleting a Captive Portal Customization	8-7
Captive Portal View	8-7
9.0 CLI Scripting.....	9-1
CLI Script.....	9-1
Pre-Configured Scripts.....	9-2
Creating a Script File	9-2
Importing a Script File	9-6
Editing a Script File.....	9-7
Sending a Script File.....	9-7
Deleting a Script File.....	9-8
Logs	9-8
Displaying a Log File	9-8
Exporting a Log File.....	9-8
Terminal	9-9
Connecting to a Device.....	9-9
Session Preferences.....	9-9
10.0 Control Panel.....	10-1
Watchdog	10-1
Scheduler Jobs.....	10-2
Starting/Stopping Scheduled Jobs.....	10-2
Viewing Scheduler Jobs	10-2
Editing a Scheduled Job	10-3
Deleting a Scheduled Job	10-4
Scheduler History	10-4
Session Management.....	10-4
Session Information.....	10-4
11.0 Discovery.....	11-1
Managed Devices.....	11-2
Discovering/Re-Discovering Devices	11-2
Adding a Device	11-3
Cloning a Device	11-6
Editing a Device.....	11-6
Deleting a Device	11-7
Searching for a Device.....	11-7
Perform Device Operations	11-7
Discovery Displays	11-8
Discovery Profiles.....	11-12
Creating a Discovery Profile	11-13
Editing a Discovery Profile.....	11-14
Deleting a Discovery Profile.....	11-14
Profile Information	11-15

Table of Contents (continued)

Third-Party Devices Support	11-16
Adding Third-Party Device Support	11-16
Editing Third-Party Device Support	11-17
Deleting Third-Party Device Support	11-17
Mibset List	11-17
Import MIBs	11-18
Importing MIBs	11-18
Hardware Inventory	11-19
Asset Information	11-19
Ports	11-19
Port Information	11-20
Link	11-20
Creating a Link	11-21
Editing a Link	11-21
Deleting a Link	11-21
Existing Links Table	11-21
Setting Frequencies	11-22
Configuring Automatic Polling	11-22
IP Failover	11-23
Switch Monitoring	11-23
12.0 Groups	12-1
MAC Groups	12-1
Creating a MAC Group	12-1
Editing a MAC Group	12-1
Deleting a MAC Group	12-2
VLAN Groups	12-2
Creating a VLAN Group	12-2
Editing a VLAN Group	12-2
Deleting a VLAN Group	12-2
Network Groups	12-2
Creating a Network Group	12-3
Editing a Network Group	12-3
Deleting a Network Group	12-3
Multicast Groups	12-3
Creating a Multicast Group	12-3
Editing a Multicast Group	12-3
Deleting a Multicast Group	12-4
Service Groups	12-4
Creating a Service Group	12-4
Editing a Service Group	12-4
Deleting a Service Group	12-4
Services	12-4
Creating a Service	12-4
Editing a Service	12-5
Deleting a Service	12-5

Table of Contents (continued)

Service Port.....	12-5
Creating a Service Port.....	12-5
Editing a Service Port.....	12-5
Deleting a Service Port.....	12-5
13.0 PIM	13-1
PIM Global Configuration	13-1
Creating a PIM Global Profile	13-1
Editing a PIM Global Profile.....	13-2
Assigning a PIM Global Profile.....	13-2
Removing a PIM Global Profile.....	13-2
Deleting a PIM Global Profile.....	13-2
PIM Interface.....	13-2
Displaying PIM Interfaces.....	13-2
Creating a PIM Interface.....	13-2
Deleting a PIM Interface.....	13-3
PIM Candidate.....	13-3
Creating a PIM Candidate Profile.....	13-3
Editing a PIM Candidate Profile.....	13-4
Deleting a PIM Candidate Profile.....	13-4
PIM Device View	13-4
14.0 License	14-1
Node Management License	14-1
VMM License	14-2
Add/Import a New License	14-2
OmniVista Licensing Options	14-2
Node Management License Options.....	14-2
VMM License Options.....	14-2
License Management.....	14-2
Node Management License.....	14-3
VMM License.....	14-3
Add or Import License	14-4
15.0 Locator.....	15-1
Locate Screen	15-1
Search Type	15-2
Browse Screen.....	15-3
Poll Screen.....	15-3
Locate	15-3
Locating a Switch	15-3
Search Results	15-3
Browse	15-6
Browse Results.....	15-7
Poll.....	15-9
Settings	15-9
General.....	15-10
Data Retention Policy	15-10
Locator Data Statistics.....	15-10

Table of Contents (continued)

16.0 Notifications	16-1
Notifications Home	16-1
Viewing Traps	16-2
Basic Trap Information	16-3
Detailed Trap Information	16-3
Acknowledging/Renouncing/Deleting Traps	16-4
Polling Devices for Traps	16-4
Trap Definition	16-4
Viewing Trap Definitions	16-4
Editing a Trap	16-5
Trap Responder	16-5
Creating a Trap Responder	16-5
Editing a Trap Responder	16-10
Deleting a Trap Responder	16-10
Viewing Trap Responders	16-10
Trap Configuration	16-10
Device Selection	16-10
Configure Traps	16-11
Summary	16-12
Settings	16-12
Notification Configuration	16-12
Trap E-Mail Configuration	16-13
17.0 PolicyView	17-1
Creating Policies for Users and Groups	17-2
Creating Policies for Resources	17-2
Creating One Touch Policies	17-2
View/Modify Policies and Policy Lists	17-2
Expert Mode	17-2
Creating Policies for Applications	17-3
QoS-Qualified Devices	17-3
Saving Changes to the Switch	17-3
Required Traps	17-3
Policy Precedence and Conflicts	17-4
Creating Policies for Users and Groups	17-4
Unified Policies	17-4
Unified Policy List	17-12
Resource Policies	17-14
One Touch Policies	17-16
One Touch Data Policies	17-16
One Touch ACL Policies	17-18
One Touch Voice Policies	17-20
View/Modify Policies and Policy List	17-23
Policies	17-23
Policy Lists	17-24
Policies by Switch	17-25

Table of Contents (continued)

Expert Mode	17-26
Creating a Custom Policy	17-26
Applying a Custom Policy to the Network	17-26
Editing a Custom Policy	17-27
Deleting a Custom Policy	17-27
Policy Information	17-27
Config for Policy	17-28
Device Selection	17-28
Set Condition	17-28
Set Action	17-34
Validity Period	17-38
Review	17-38
Applications	17-38
Application Visibility Policies	17-38
Application Visibility Policy Lists	17-45
Policy List Information	17-46
Creating a Policy List	17-46
Editing a Policy List	17-46
Deleting a Policy List	17-46
18.0 Preferences	18-1
Locale	18-2
Theme	18-2
Inactivity Timeout	18-3
Table/List View Mode	18-3
Temperature Unit	18-3
Device Naming Pattern	18-4
Network Status Color Preferences	18-4
Alarms Color Preferences	18-4
Quarantine Manager Color Preferences	18-4
ProActive Lifecycle Management Color Preferences	18-4
Branding	18-4
Proxy	18-5
ProActive Lifecycle Management	18-5
Videos	18-5
Email	18-6
Install Zulu CEK	18-6
ProActive Lifecycle Management Overview	18-6
Setting Up ProActive Lifecycle Management	18-6
Viewing Information	18-8
Manually Uploading Information	18-9

Table of Contents (continued)

19.0 Quarantine Manager	19-1
Quarantine Manager Requirements	19-2
Hardware/Software Requirements	19-2
Candidates	19-4
Candidates Quarantine List	19-4
Banned	19-5
Adding a Device to the Banned List	19-5
Editing a Device on the Banned List	19-5
Releasing a Device from the Banned List	19-5
Retry	19-5
Re-Polling for Banned Devices	19-6
Banned Quarantine List	19-6
Never Banned	19-6
Adding a Device to the Never Banned List	19-6
Editing a Device on the Never Banned List	19-6
Deleting a Device from the Never Banned List	19-7
Never Banned Quarantine List	19-7
Disabled Ports	19-7
Release a Device from the Disabled Port List	19-7
Edit a Device in the Disabled Port List	19-7
Retry a Port Operation	19-7
Disabled Port List	19-8
Rules	19-8
Quarantine Manager Rule Overview	19-8
Creating a Quarantine Manager Rule	19-12
Editing a Quarantine Manager Rule	19-13
Deleting a Quarantine Manager Rule	19-13
Enabling/Disabling a Quarantine Manager Rule	19-13
Importing a Quarantine Manager Rule	19-13
Quarantine Manager Rule List	19-14
Configuration	19-15
Quarantined Manager Remediation (QMR)	19-15
Configuring Quarantine Manager	19-15
Configuring Quarantine Manager on OmniAccess WLAN Devices	19-17
Responders	19-17
Creating a Quarantine Manager Responder	19-18
Editing a Quarantine Manager Responder	19-18
Deleting a Quarantine Manager Responder	19-19
Automatic Event Responders List	19-19
TAD Profile	19-19
Creating a TAD Profile	19-19
Assigning a TAD Profile	19-20
Editing a TAD Profile	19-21
Deleting a TAD Profile	19-21
Monitoring Group List	19-21

Table of Contents (continued)

TAD View	19-21
Monitoring Groups	19-22
Port Ranges.....	19-22
Statistics Port.....	19-22
Statistics Anomaly Traffic	19-22
Statistics Anomaly Summary	19-23
Settings	19-23
SysLog Listener.....	19-23
SysLog Generator Target	19-23
20.0 Report	20-1
Report Configuration	20-1
Creating a Report	20-2
Editing a Report.....	20-2
Deleting a Report.....	20-3
Report List.....	20-3
21.0 Resource Manager.....	21-1
Backup	21-1
Performing a Backup.....	21-2
Editing a Configuration Backup File	21-4
Deleting a Backup	21-4
Backup Information.....	21-4
Important Facts About Backing Up.....	21-5
Restore.....	21-5
Performing a Restore.....	21-6
Compare.....	21-7
Selecting Files	21-7
Comparing Files	21-7
Backup Summary	21-7
Upgrade Image.....	21-8
Importing the Upgrade Files	21-8
Installing the Upgrade Files	21-8
ISSU Upgrade	21-10
Important Information About Upgrades	21-11
If FTP User Names/Passwords are Undefined	21-11
If Version Numbers are Older	21-12
File Sets Information.....	21-12
Inventory.....	21-12
Creating an Inventory Report.....	21-12
Auto Configuration.....	21-13
Auto Configuration Overview	21-13
Auto Configuration on a New Switch	21-14
Quick Steps for Automatic Remote Configuration	21-15
Creating an Instruction File.....	21-15
Editing an Instruction File.....	21-16
Deleting an Instruction File	21-17
The Instruction Files List.....	21-17

Table of Contents (continued)

Switch File Set.....	21-18
Overview.....	21-18
Creating a Switch File Set.....	21-19
Assigning a Switch File Set	21-19
Editing a Switch File Set	21-19
Deleting a Switch File Set.....	21-19
Settings	21-20
Backup Retention Policy.....	21-20
22.0 SIP	22-1
SIP Overview.....	22-2
Active Calls.....	22-2
Viewing Active Call Records.....	22-3
Graphical View	22-4
Ended Calls.....	22-5
Viewing Ended Call Records	22-5
Aggregated Records.....	22-5
Detailed Records	22-5
One Touch Profile	22-7
Creating a SIP One Touch Profile	22-8
Editing a SIP One Touch Profile.....	22-8
Deleting a SIP One Touch Profile.....	22-8
Removing a SIP One Touch Profile.....	22-9
Applying a SIP One Touch Profile	22-9
Viewing SIP One Touch Profiles.....	22-9
SIP Profile	22-9
Creating a SIP Profile	22-9
Editing a SIP Profile.....	22-10
Deleting a SIP Profile.....	22-10
Applying a SIP Profile	22-10
Viewing SIP Profiles	22-11
Global Params Profile.....	22-12
Trusted Servers Profile.....	22-13
Threshold Profile	22-14
SOS Profile.....	22-15
TCP Port Profile.....	22-16
UDP Port Profile	22-17
Device View.....	22-18
SIP Profile.....	22-19
TCP Port Profile.....	22-19
UDP Port Profile	22-19
Global Param Profile.....	22-19
Trusted Server Profile.....	22-20
Threshold Profile	22-20
SOS Profile.....	22-20
SIP Statistics	22-20
SIP Settings.....	22-21

Table of Contents (continued)

23.0 Topology	23-1
Topology Maps.....	23-2
Working with Topology Maps.....	23-2
Working with Network Devices	23-9
24.0 Unified Access	24-1
Unified Profile.....	24-1
Authentication and Classification.....	24-2
Configuring Unified Profile	24-3
Device Config	24-28
Profile Polling.....	24-44
Unified Policies.....	24-44
Creating Unified Policy.....	24-45
Applying a Unified Policy to the Network	24-45
Editing a Unified Policy	24-45
Deleting Unified Policy	24-45
Policy Information	24-46
Unified Policy List	24-52
Unified Policy List Information	24-52
Creating a Unified Policy List.....	24-53
Applying a Unified Policy List to the Network.....	24-54
Editing a Unified Policy List	24-54
Deleting a Unified Policy List	24-54
Unified Policy List Information	24-54
mDNS.....	24-55
mDNS Flow	24-55
mDNS Configuration Screen	24-56
Configuring mDNS for a Device.....	24-56
Editing an mDNS Configuration.....	24-57
Deleting an mDNS Configuration.....	24-57
Viewing mDNS Configurations	24-57
Premium Services	24-58
BYOD/ClearPass Overview	24-58
BYOD/ClearPass Setup.....	24-60
ClearPass	24-63
Consistency Check.....	24-66
Performing a Consistency Check	24-66
Authentication Records.....	24-66
25.0 Users and User Groups	25-1
Security Levels	25-2
Default Groups, Users, Roles.....	25-2
Working with User Groups, Users, and User Roles.....	25-3
Role Management.....	25-3
Creating a User Role	25-3
Editing a User Role.....	25-4
Deleting a User Role.....	25-4
Existing Roles Table.....	25-5
User Role Feature	25-5

Table of Contents (continued)

Group Management	25-6
Creating a User Group.....	25-6
Editing a User Group	25-6
Deleting a User Group.....	25-6
Existing Groups Table	25-7
User Management.....	25-7
Creating a User	25-7
Editing a User	25-7
Deleting a User.....	25-7
Existing Users Table.....	25-8
Authentication Server	25-8
26.0 VLAN Manager	26-1
VLAN Overview	26-1
VLANs Table	26-2
Creating a VLAN.....	26-3
Copying a VLAN.....	26-3
Deleting a VLAN.....	26-4
VLAN Actions	26-4
VLAN Details	26-4
MVRP	26-7
IP Interface.....	26-8
Poll	26-8
Spanning Tree.....	26-8
Summary View	26-8
Bridge View	26-9
Port View	26-10
Editing an STP Configuration	26-11
IP Router	26-12
Creating an IP Router Interface.....	26-12
Editing an IP Router Interface.....	26-13
Deleting an IP Router Interface	26-13
Viewing IP Router Interfaces	26-13
VLAN Details.....	26-14
Device View.....	26-14
Port View	26-15
Link Aggregate View.....	26-15
Device Selection.....	26-16
VLAN Configuration.....	26-17
Default Ports Assignment.....	26-17
Q Tagged Ports Assignment	26-17
Review	26-17
MVRP Overview	26-18

Table of Contents (continued)

Summary View	26-18
Switches List.....	26-18
Ports Configuration	26-19
Global Parameters.....	26-19
Port Parameters	26-20
Devices/Ports	26-21
Review.....	26-21
IP Interface.....	26-21
Creating an IP Interface.....	26-21
Editing an IP Interface	26-22
Deleting an IP Interface	26-22
Viewing IP Interfaces.....	26-22
Poll	26-24
Multiple Virtual Routing and Forwarding.....	26-24
Configuring the Multiple VRF Feature.....	26-25
27.0 VM Manager.....	27-1
Virtualization/VM Manager Overview	27-2
Virtualization	27-2
VM Manager and Virtualization.....	27-2
VM Manager and Traffic Shaping	27-5
Configuring VM Manager.....	27-5
VM Manager Configuration Quick Steps.....	27-5
Hypervisor Systems	27-6
Creating a VM Server	27-6
Editing a VM Server.....	27-7
Deleting a VM Server.....	27-7
Hypervisor Systems Table.....	27-7
VM Locator - Host Networks.....	27-8
Searching for Host Machines.....	27-8
VM Locator - VM Networks.....	27-9
Searching for Virtual Machines.....	27-9
VM Networks Table	27-10
Exclude VLAN	27-11
Creating an Exclude VLAN	27-11
Editing and Exclude VLAN.....	27-11
Deleting an Exclude VLAN	27-11
Exclude VLAN List.....	27-11
VM VLAN Configuration	27-12
VM VLAN Configuration - Apply UNP VLAN.....	27-12
VM VLAN Configuration - Enable MVRP Ports.....	27-12
VM VLAN Configuration - Enable One-Touch SPB	27-13
VLAN Notification	27-13
Resolving a VM Configuration Problem.....	27-13
Notification List	27-14
VMM Devices List.....	27-14
VMM Devices List.....	27-14

Table of Contents (continued)

Settings - VM Polling	27-16
Settings - SBP	27-16
28.0 VXLANs.....	28-1
VXLAN Service.....	28-1
Creating a VXLAN Service.....	28-2
Creating a VXLAN Network Profile	28-2
Creating an Service Distribution Point (SDP) Tunnel	28-2
Re-Applying a VXLAN Service	28-3
Editing a VXLAN Service	28-3
Deleting a VXLAN Service.....	28-3
SAP Profile	28-3
Creating a SAP Profile	28-4
Assigning a SAP Profile	28-4
Editing a SAP Profile	28-4
Deleting a SAP Profile	28-5
Access Port Profile	28-5
Creating an Access Port Profile.....	28-5
Editing an Access Port Profile	28-5
Deleting an Access Port Profile	28-5
VXLAN Device View	28-5
VXLAN Information.....	28-5
VM Snooping Overview	28-7
VM Snooping Profile	28-7
Editing a VM Snooping Profile	28-8
Deleting a VM Snooping Profile.....	28-9
Removing a VM Snooping Profile From a Switch/Port.....	28-9
VM Snooping Statistics.....	28-9
Device View	28-10

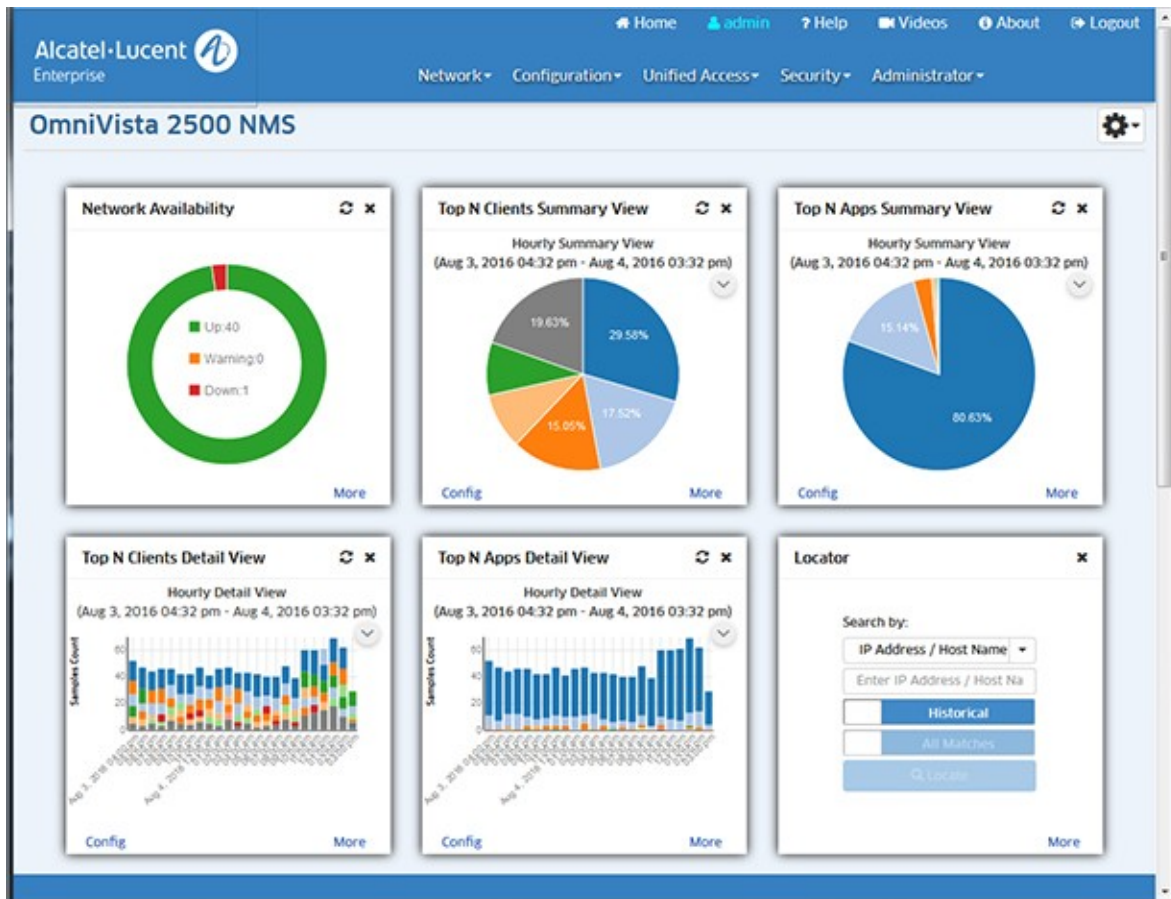
1.0 Getting Started with OmniVista 2500 NMS

Alcatel-Lucent Enterprise's OmniVista 2500 NMS uses a web-based user interface. The Web GUI is supported on the following browsers: Internet Explorer 10+ (on Windows client PCs), Firefox 26+ (on Windows and Redhat/SuSE Linux client PCs), Chrome 26+ (on Windows and Redhat/SuSE Linux client PCs). This Quick Start guide will get you up and running with OmniVista 2500 NMS. The sections below provide an overview of the OmniVista 2500 NMS.

- OmniVista 2500 NMS User Interface
 - Applications
 - Customizing the Dashboard
 - Configuration/Display Buttons
 - Working with Tables
- Setting Up OmniVista 2500 NMS for Network Management
 - Discovering Devices
 - Editing Discovered Devices
 - Configuring Traps
 - Saving Changes

OmniVista 2500 NMS User Interface

The OmniVista 2500 NMS Home Page displays a dashboard with application widgets that provide an overview of key applications. The widgets also link directly to their application, so you can access the application for more detailed information/configuration. The dashboard is customizable. You can add/remove widgets, and drag and drop them into any order, as well as close widgets and refresh widget content.



The following links are also available on the Home Page and at the top of every screen in OmniVista 2500 NMS:

- **Home** - Returns the user to the Dashboard.
- **Admin** - Brings up the Local User Management Screen.
- **Help** - Brings up the OmniVista 2500 NMS Getting Started Guide.
- **Videos** - Launches the Alcatel-Lucent Enterprise YouTube Demo Playlist.
- **About** - Displays basic OmniVista 2500 NMS information (e.g., Build Number and date).
- **Logout** - Logs you out of OmniVista 2500 NMS.

Applications

Applications in OmniVista 2500 NMS are organized as shown below.

Network

- Discovery
- Topology
- Locator
- Notifications
- VM Manager (VM Manager License required)
- Analytics
- Application Visibility

Configuration

- VLANs
- VxLANs
- IP Multicast
- CLI Scripting
- PolicyView
- SIP
- Captive Portal
- Groups
- Application Launch
- Report
- Resource Manager

Unified Access

- Unified Profile
- Unified Policy
- Multimedia Services (mDNS)
- Premium Services (BYOD)

Security

- Users and User Groups
- Authentication Servers
- Quarantine Manager


Administrator

- Control Panel
- Preferences
- Audit
- License

Customizing the Dashboard


You arrange widgets on the Dashboard by clicking on a widget and dragging it to a new location. You can also add or remove widgets from the Dashboard and customize the Dashboard layout using the Dashboard Settings option. This section provides an overview of application widgets and the Dashboard. For detailed information see the Dashboard online help.

Adding Widgets

To add a widget to the Dashboard, click on the Settings icon  and select **Add Widget**. The Add Widget Screen, shown below is displayed. Widgets that are already displayed on the Dashboard are displayed in grey. Available widgets are displayed in bold. You can select a widget to display a sample of the widget and a brief summary of the information provided by the widget.

Select an available widget from the list and click **OK**. The widget will be added to the upper-left of the dashboard. You can then drag and drop the widget to any position in the Dashboard. Note that you can only add one widget at a time. Repeat the procedure to add additional widgets

Removing Widgets

To remove a widget from the dashboard, click on the Delete icon (**x**) at the top-right corner of the widget. You can add the widget back or add a new widget by clicking on the Settings icon  and selecting **Add Widget**.

Additional Options Provided Within Widgets

You can configure Dashboard display settings for certain widgets (e.g., Analytics Statistics) by clicking on the **Config** link at the bottom of the widget. You can then modify certain display options, such as the number of items displayed in the widget, the time period displayed in the widget, etc. This modifies the display on the Dashboard only. It does not modify the display within the application.

You can also hover over certain displays (e.g., Analytics) to view more detailed information, and click on data within certain widgets to display additional information (e.g., clicking on bar and donut graphs displays additional information such as numeric values, clicking on Quarantine Manager topics at the bottom of the graph displays additional information for each topic).

Click on the **More** link at the bottom of a widget to open the full application for the widget. To return to the Dashboard, click on the Alcatel-Lucent Enterprise logo at the top left corner of the screen.

The Favorites Widget and Adding Items to the Favorites Tab


The Favorites widget allows users to create a list of “quick links” for easy access to any OmniVista 2500 NMS application. To add an application to the Favorites list, start typing the name of the application in the text box. As text is entered, auto-complete displays a list of matching applications. Click an application to select it, then click the add icon (+).

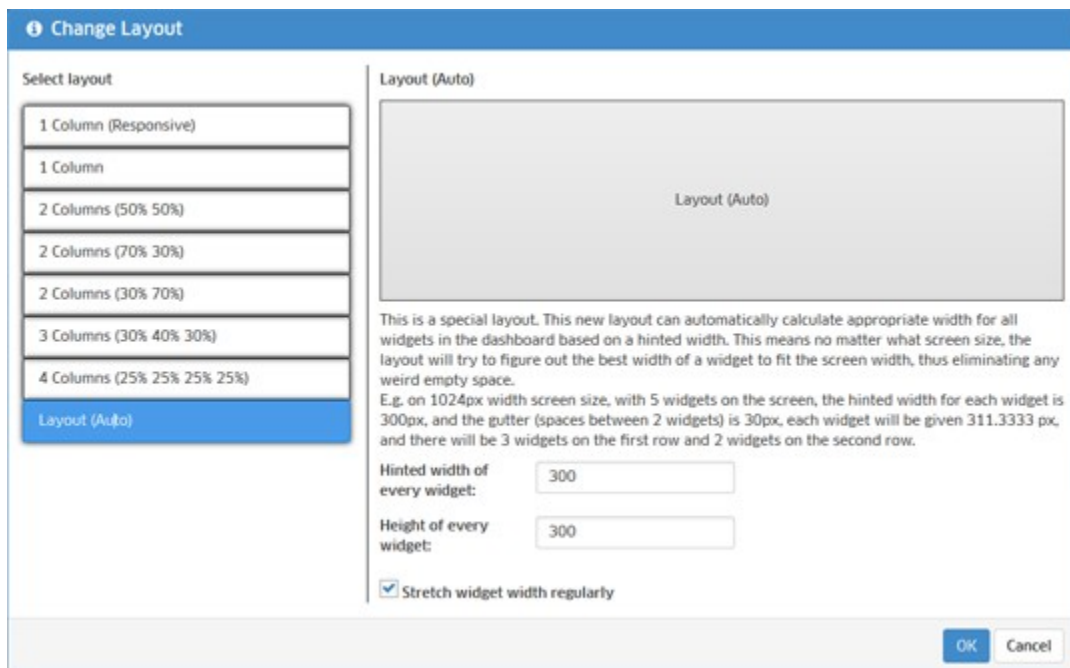
To replace an application in the list of Favorites, click the edit icon. Begin typing the new application name. Select the application and click the Save icon.

Note: Once an application has been added to the Favorites widget, a link to the application will also display under the Favorites tab in the OmniVista 2500 NMS Main Navigation Menu.

To remove an application from the Favorites widget, click the Delete icon (x) next to the application link.

Customizing the Dashboard Layout

The Settings option also enables you to customize the Dashboard layout. To customize the layout, click on Settings icon  and select **Change Layout**. The Change Layout Screen appears. By default, widgets are displayed in the "Auto" layout. To change the layout, select an option on the left. When you select a layout on the left, the layout is displayed in the Layout area. To choose that layout, click **OK**. The change takes effect immediately. If necessary, you can re-arrange the widgets in the new layout, by clicking and dragging them to new locations on the Dashboard.



Configuration/Display Buttons

OmniVista 2500 NMS provides standard tools for interacting with configuration/display screens. These buttons include:















Configuration Buttons	
	Create/Add Click the Create/Add button to create a new entry within the configuration screen.
	Edit To edit an existing entry, select the entry in the configuration screen and click the Edit button.
	Copy/Duplicate Used to duplicate an existing entry in a list or table. Select an entry in a table, and click on the Copy/Duplicate button to bring up a configuration screen with the selected configuration. You can then modify certain fields to quickly create a new entry. This can save time when creating multiple entries with similar configurations.
	Delete To delete an entry, select the entry and click the Delete button.

Table Buttons	
	<p>Table View Table view organizes configuration screens in a table-based layout that is viewed and sorted similar to a spreadsheet. Table data may be exported to CSV files.</p>
	<p>List View List View organizes information and configuration tasks in a list view. Note that information displayed in List View cannot be Printed via the OmniVista 2500 NMS Print button. To print information from a particular screen, switch to Table View (if available).</p>
	<p>Settings Used to configure the column headings to display in a table, Click on the Settings button and select the column headings you want to display.</p>
	<p>Refresh The Refresh button loads the latest data for an application table, chart or list.</p>
	<p>Help Click on the Help button to load for a context-sensitive walk-through of an application, table or configuration screen.</p>
	<p>Sort Information displayed in List View may be sorted in alphabetical order, either ascending or descending, by clicking on the Sort button. You can also click on the Up/Down arrows at the top of any table column in Table View to sort the data in ascending or descending order based on the selected column.</p>
	<p>Search Click the Search button and enter search criteria in the "Search..." field to display specific entries in the table. Click on the "x" to the right to return to the original display.</p>
	<p>Filter Users can create custom filters for OmniVista 2500 NMS tables to display specific data. To create a custom filter:</p> <ul style="list-style-type: none"> • Click the filter button. • In the Filter screen, click Add. • Enter a Filter Name for the custom filter. Enter a Filter Description for the filter. • In the Conditions area, for "Filter elements where ... of the following apply", specify the strictness of the conditions (e.g., display results if ANY condition is encountered, or only display results if ALL of the conditions are encountered). • Select an option from the application-specific conditions list. This automatically enables the conditions. • Next, fine-tune conditions by selecting options under the following pull down menus: <ul style="list-style-type: none"> ○ have/not have ○ Name - select a column heading to use in the filter. ○ contains <ul style="list-style-type: none"> ○ begins with ○ ends with ○ equal ○ not equal ○ [enter value] - enter a specific text string ○ non-case sensitive/case sensitive • Additional application-specific conditions may be added to the conditions outlined in the previous step by clicking the "Add new Condition" link. • You can also add a completely new group of conditions by clicking the Add icon and repeating the steps outlined above.
	<p>Reset Click the Reset button after filtering a table to return to the original display.</p>
	<p>Export to CSV Click the CSV button to download information displayed in Table View to a CSV (spreadsheet) file.</p>

<p>Add to Report</p>	<p>Add To Report Click to create a report for the page. These reports are PDF versions of tables and reports generated in certain OmniVista applications (e.g., Discovery, Locator, Analytics). Basically, in addition to viewing information in real-time in OmniVista (e.g., Discovery Inventory List, Analytics Utilization Reports), you can generate PDFs of the screens. When a report is generated, it takes a current snapshot of the application information. These reports can be generated immediately or you can schedule them to be generated at regular times/intervals (e.g., Daily, Weekly). You can also configure a report to be e-mailed when it is generated. Reports are configured in the Report application.</p>
<p>Print</p>	<p>Print Print Table View information by clicking the Print button. Note that information displayed in List View cannot be Printed via the OmniVista 2500 NMS Print button (button is not available). To print information from a screen, switch to Table View (if available).</p>

Working with Tables

Information in OmniVista is primarily presented in table format. There are common functions/behaviors for tables in OmniVista. the general functionality of each area is described below. Details for each button are provided in the Configuration/Display Buttons section.

The screenshot shows the 'CLI Scripting' interface. At the top, there are navigation tabs: Home, Configuration, Unified Access, Security, and Administrator. The main content area is titled 'CLI Script' and contains a 'CLI Script List' table. The table has columns for 'Filename', 'Shared Admin Script', 'Description', and 'Last Modified'. Below the table, there are controls for 'Show: All' and 'Showing All 5 rows'. On the right side of the interface, there are several icons for configuration and display options, including a search icon, a filter icon, an export icon, and a print icon. Labels with arrows point to these icons and the table controls, explaining their functions.

Filename	Shared Admin Script	Description	Last Modified
shadmin_Ewebimageclean.sc...	true	Clean image file from Workl...	Aug 30, 2016
shadmin_save_config_snapsh...	true	Save to Certified and Workin...	Aug 30, 2016
shadmin_sFlowTool_config.sc...	true	Sflow - Setup sampler and p...	Aug 30, 2016
shadmin_show_switch.script	true	Show Switch Information - A...	Aug 30, 2016
shadmin_spantree_view.script	true	<N/A>	Aug 30, 2016

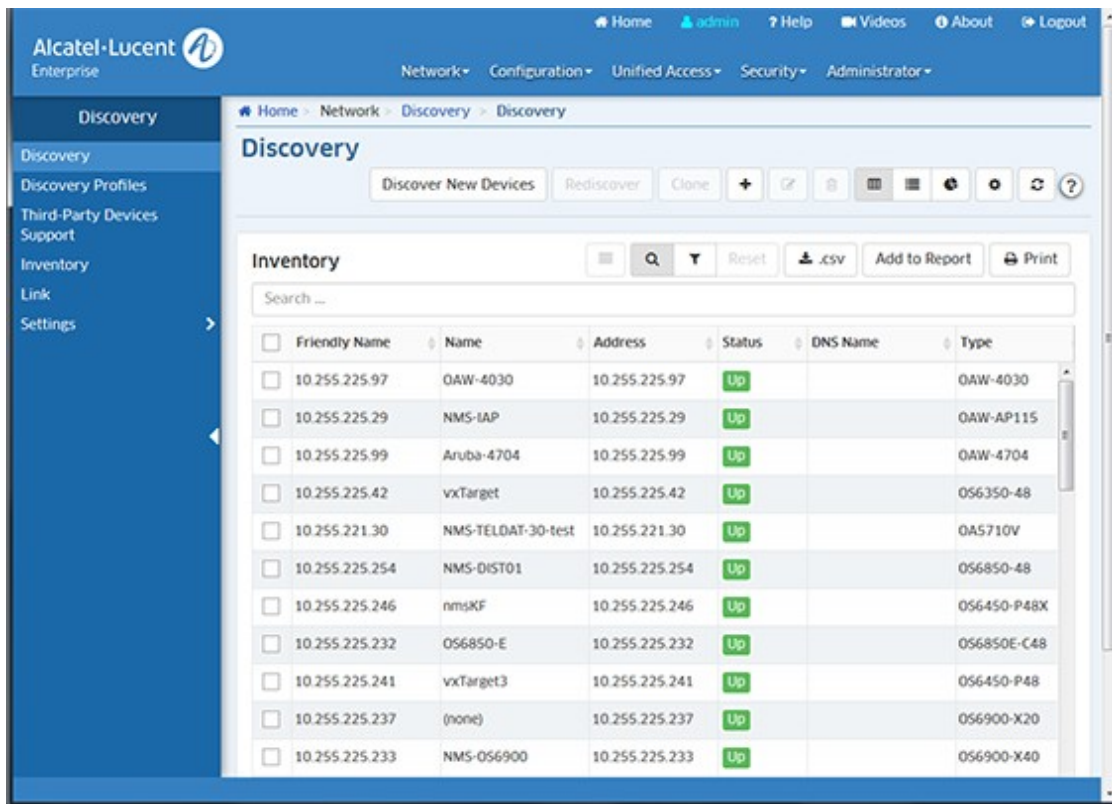
- **Configuration Options** - Used to create, edit, delete entries (e.g., create, edit, delete a CLI Script). Details for each button are provided in the Configuration/Display Buttons section.
- **Display Options** - Used to change the table display from Table View to List view, to set the columns you want to display, and to refresh the data in the table. Details for each button are provided in the Configuration/Display Buttons section.
- **Filter, Export, Print** - You can filter the data that is displayed by clicking on the Filter icon and selecting/creating a filter. You can also export the table into a .csv file, or print the table. Details for each button are provided in the Configuration/Display Buttons section.

- **Search** - Enter search criteria in the **Search...** field to display specific entries in a table. As you enter criteria, only those entries matching the criteria will appear in the table. Click on the **Reset** button to return to the original table display.
- **Sort** - Click on one of the arrows at the top of a column to sort the table in ascending or descending order based on the column.
- **Set Lines to Display/Page Through Table** - Set the number of lines to display in the table using the **Show** drop-down menu at the bottom left corner of the table. You can also page through a large table using the arrows at the bottom right corner of the table.

Setting Up OmniVista 2500 NMS for Network Management

Discovering Devices

After you log into OmniVista, open the Discovery application to discover devices on your network. Click on the **Discover New Devices** button and enter a range of devices to discover. You can use the Default Discovery Profile for your initial discovery. (A Discovery Profile contains the parameters that are used by OmniVista when performing a discovery (e.g., SNMP version used to discover devices, FTP/Telnet passwords needed to connect to a device). When the discovery is complete, all devices discovered display in the Inventory List on the Discovery Screen.



Friendly Name	Name	Address	Status	DNS Name	Type
<input type="checkbox"/>	10.255.225.97	OAW-4030	10.255.225.97	Up	OAW-4030
<input type="checkbox"/>	10.255.225.29	NMS-IAP	10.255.225.29	Up	OAW-AP115
<input type="checkbox"/>	10.255.225.99	Aruba-4704	10.255.225.99	Up	OAW-4704
<input type="checkbox"/>	10.255.225.42	vxTarget	10.255.225.42	Up	056350-48
<input type="checkbox"/>	10.255.221.30	NMS-TELOAT-30-test	10.255.221.30	Up	0A5710V
<input type="checkbox"/>	10.255.225.254	NMS-DIST01	10.255.225.254	Up	056850-48
<input type="checkbox"/>	10.255.225.246	nmsKF	10.255.225.246	Up	056450-P48X
<input type="checkbox"/>	10.255.225.232	056850-E	10.255.225.232	Up	056850E-C48
<input type="checkbox"/>	10.255.225.241	vxTarget3	10.255.225.241	Up	056450-P48
<input type="checkbox"/>	10.255.225.237	(none)	10.255.225.237	Up	056900-X20
<input type="checkbox"/>	10.255.225.233	NMS-056900	10.255.225.233	Up	056900-X40

Editing Discovered Devices

Once switches are discovered, you may want or need to edit entries in the Inventory List. To edit a discovered device, select the device in the Inventory List and click on the Edit icon to display the Edit Discovery Manager Entry window, shown below.

Why Edit a Discovered Device?

To Redefine the Primary IP Address

When switches are auto-discovered via a Ping Sweep or ARP discovery, each IP address in a range or subnet is pinged. OmniVista uses the first IP address that responds to a ping as that device's primary IP address. However, if multiple VLANs exist in the device, additional IP addresses in the device will also respond to pings. The Edit Discovery Manager Entry window's **IP Address** field combo box lists these additional IP addresses and enables you to select any address listed as the device's primary IP address.

To Specify the Correct Write Community Name

All devices that are discovered are initially specified to have the default write community name, **public**. If any discovered devices in your network have a non-default write community name, use the Edit Discovery Manager Entry window's SNMP Settings tab to specify the correct community name to OmniVista. If the correct write community name is not specified to OmniVista, you will not be able to write configuration changes to the switch.

Please Note: Switches' SNMP write (set) community names are not configurable from OmniVista. SNMP read (get) and write (set) community names can only be configured by logging onto the switch.

To Specify the CLI or FTP User Name and Password

Firmware configuration files for AOS devices can be saved to the OmniVista server and restored when desired. When files are saved, they are FTPed from the switch to the OmniVista server. When files are restored, they are FTPed from the server to the switch. New configuration files can also be installed via FTP. To FTP files, OmniVista must know the FTP login name and password that is defined on the switch. The **CLI/FTP User Name** and **CLI/FTP Password** fields on the Edit Discovery Manager Entry window enable you to specify this information to OmniVista.

Please Note:

- If you do not define the FTP login names and passwords and you attempt to save, restore, or update configuration files for these devices, you will be individually queried for the FTP login name and password of each individual switch for which configuration files are being saved, restored, or updated.
- The user names and passwords entered in these fields are used for FTP ONLY. They are not used for Telnet. When you Telnet to a device, you will be queried for a user name and password.

To Redefine the SNMP Version or SNMP Parameters

The Edit Discovery Manager Entry window enables you to redefine the SNMP version that OmniVista uses to communicate with AOS devices. You can also redefine SNMP parameters. The SNMP version or parameter settings that OmniVista uses cannot be changed until OmniVista has connected to the switch. AOS devices support SNMP version 1, SNMP version 2 or SNMP version 3.

Configuring Traps

It is necessary to configure the switches in the network to send OmniVista the traps that are needed by different applications. Traps are configured in the Notifications application. The traps OmniVista needs for each application are listed below.

Traps Needed for Topology

coldStart, warmStart, linkUp, linkDown

Note: For proper link display in Topology, linkUp and linkDown traps must be enabled for each individual port.

Traps Needed for PolicyView QoS

policyEventNotification

Saving Changes

The directory structure that stores AOS image and configuration files in flash memory is divided into two parts:

- The certified directory contains files that have been certified by an authorized user as the default configuration files for the switch. When the switch reboots, it will automatically load its configuration files from the certified directory if the switch detects a difference between the certified directory and the working directory. (Note that you can specifically command a switch to load from either directory.)
- The working directory contains files that may -- or may not -- have been altered from those in the certified directory. The working directory is a holding place for new files to be tested before committing the files to the certified directory. You can save configuration changes to the working directory. You cannot save configuration changes directly to the certified directory.

Note that the files in the certified directory and in the working directory may be different from the running configuration of the switch, which is contained in RAM memory. The running configuration is the current operating parameters of the switch, which are originally loaded from the certified or working directory. Modifications made to the running configuration of the switch must be saved to the working directory (or lost). The working directory can then be copied to the certified directory if and when desired.

When changes are made to the configuration of an AOS device -- such as configuring the traps the switch should transmit -- the change is written to the running configuration of the switch. However, if the switch is powered off, the running configuration will be lost. To make changes to the running configuration persistent, you must save the running configuration to the working directory of the switch. You should also then copy the working directory to the certified directory, so the changes will be persistent when the switch is booted from the certified directory.

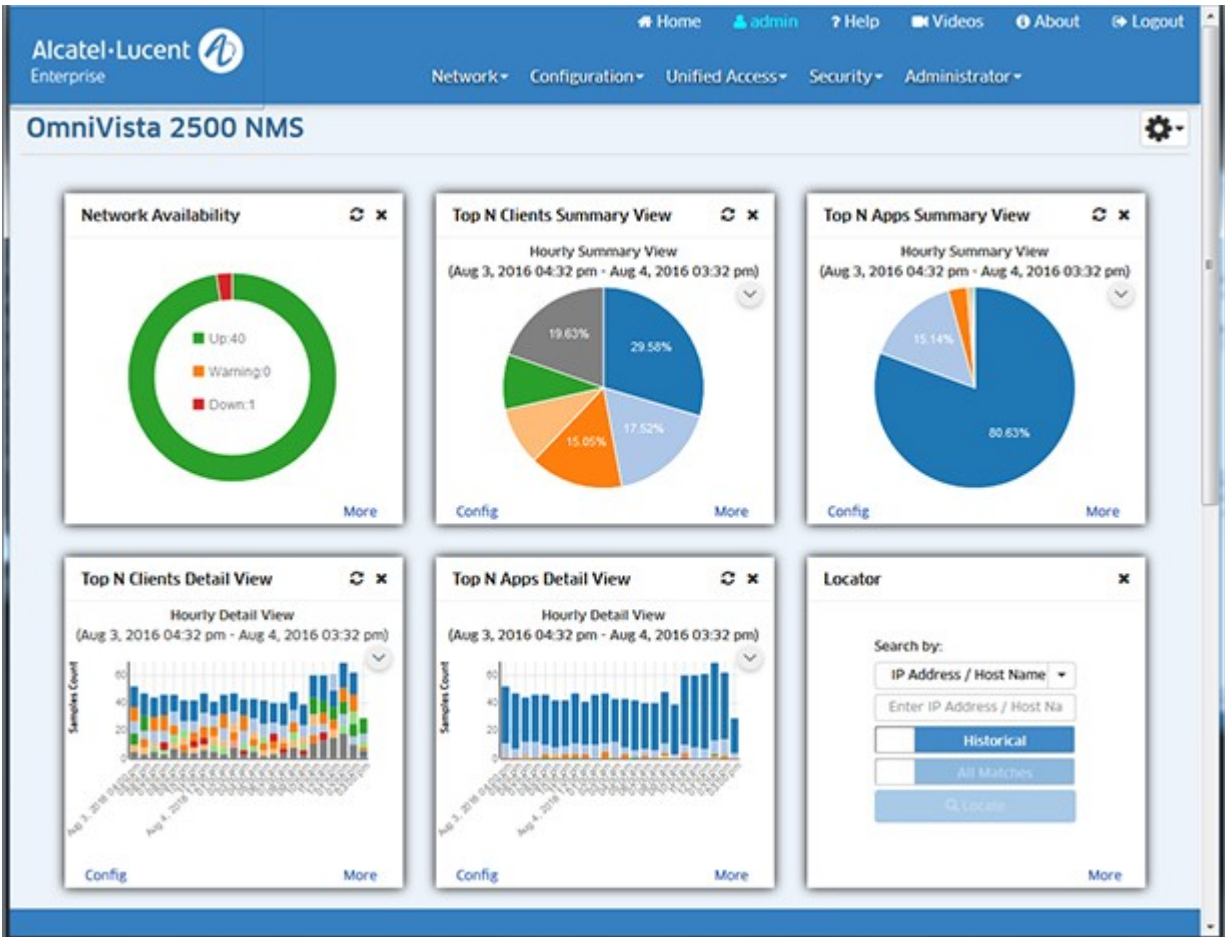
You can perform these operations in the Topology application by selecting a device(s) and clicking on the application operation (Copy Working/Running to Certified).

PolicyView QoS

When PolicyView QoS is executed, it writes the address of the LDAP server to each QoS-enabled switch in the Discovery Inventory List. (All AOS devices are QoS-enabled.) In the case of AOS devices, the LDAP address is written to the running configuration of the switch. For this reason, once PolicyView QoS has executed, all AOS devices will be left with their running configuration in the **Unsaved** state (indicating that the running configuration has changes that have not been saved to the working directory). It is important to save the running configuration to the working directory and then the certified directory after PolicyView QoS has executed. To do this, follow the steps in Saving Changes, above.

2.0 Dashboard

The OmniVista 2500 NMS Home Page displays a dashboard with application widgets that provide a quick overview of key applications. The widgets also link directly to their application, so you can access the application for more detailed information/configuration. The Dashboard is customizable; you can add/remove widgets, and drag and drop them into any order on the Dashboard.




The following links are available on the Home Page and at the top of every screen in OmniVista 2500 NMS:

- **Home** - Returns the user to the Dashboard.
- **Admin** - Brings up the Local User Management Screen.
- **Help** - Brings up the OmniVista 2500 NMS Getting Started Guide.
- **Videos** - Launches the Alcatel-Lucent Enterprise YouTube Demo Playlist.
- **About** - Displays basic OmniVista 2500 NMS information (e.g., Build Number and date).
- **Logout** - Logs you out of OmniVista 2500 NMS.

Overview

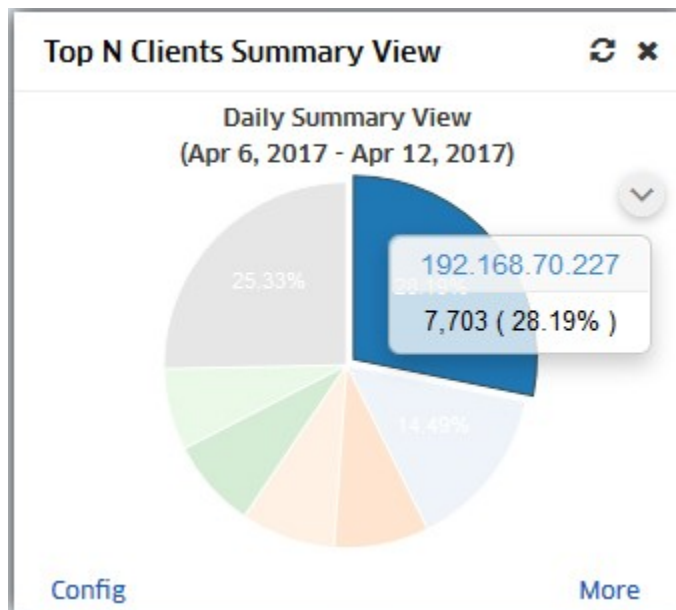
Application widgets provide a quick overview of key applications. You can hover the mouse over a widget or click on an area of the widget to display additional information, display a legend for the information in the widget, immediately refresh widget information, remove a widget from the Dashboard, configure the widget display, and directly access an application through a widget for detailed information or application configuration. Note that the options available (e.g., display addition information, configure the widget display) vary depending on the widget.



You can also customize the Dashboard by arranging widgets on the Dashboard, adding/removing widgets to/from from the Dashboard, and changing the Dashboard layout.

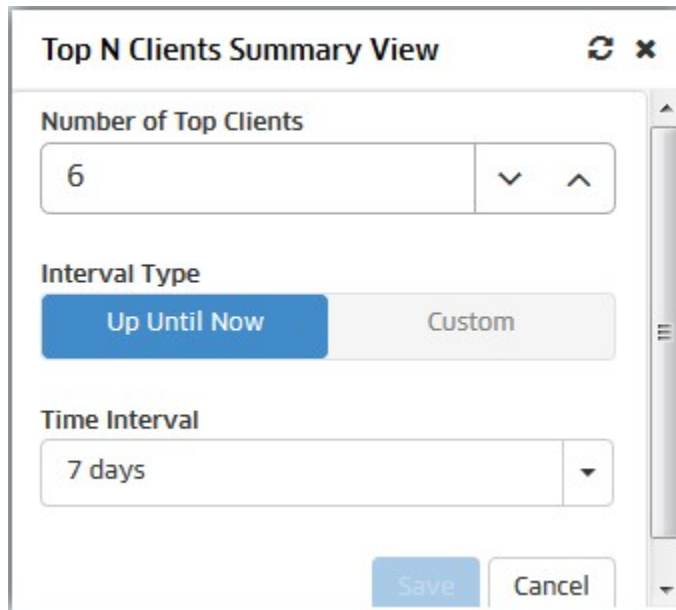
Viewing Additional Information

You can hover over certain widgets to display additional information, or click on an area for more detailed information. For example, the Top N Clients Summary Widget displays information in a pie chart format for the top network clients as a percentage of the total number traffic flows. Hover the mouse over an area of the chart to view the client IP address and the total number of flows for the client. In the example below client 192.168.70.227 is generating 7,703 flows, or 28.19% of the total flows. You can also click on an area of the chart to view detailed information about the types of traffic flows generated by the client. Information for the selected client is displayed. To return to the original view, click the Back arrow (<) in the widget.



Configuring the Widget Display

You can configure Dashboard display settings for certain widgets (e.g., Top N Clients Summary View) by clicking on the **Config** link at the bottom of the widget. You can then modify certain display options, such as the number of items displayed in the widget, the time period displayed in the widget, etc. Make any modifications and click on the **Save** button. The widget will be displayed with the new configuration.



The screenshot shows a configuration dialog box titled "Top N Clients Summary View". It contains three main sections: "Number of Top Clients" with a text input field containing "6" and up/down arrows; "Interval Type" with two radio buttons, "Up Until Now" (selected) and "Custom"; and "Time Interval" with a dropdown menu showing "7 days". At the bottom right, there are "Save" and "Cancel" buttons. The dialog box has a standard window title bar with a refresh icon and a close button.

Note: This modifies the display on the Dashboard only. It does not modify the display within the application.

Linking to the Application


Click on the **More** link at the bottom of a widget to go to the page in the full application. For example, if you clicked on the link for the Top N Clients Summary View widget, you would go to Top N Clients Report page in the application, shown below. To return to the Dashboard, click on the Alcatel-Lucent Enterprise logo at the top left corner of the screen.



Customizing the Dashboard

You arrange widgets on the Dashboard by clicking on a widget and dragging it to a new location. You can also add/remove widgets to/from the Dashboard, and customize the Dashboard layout.

Adding Widgets

To add a widget to the Dashboard, click on the Settings icon  and select **Add Widget**. The Add Widget Screen, shown below is displayed. Widgets that are already displayed on the Dashboard are displayed in gray. Available widgets are displayed in bold. You can select a widget to display a sample of the widget and a brief summary of the information provided by the widget.

Select an available widget from the list and click **OK**. The widget will be added to the upper-left of the dashboard. You can then drag and drop the widget to any position in the Dashboard. Note that you can only add one widget at a time. Repeat the procedure to add additional widgets.

All widgets are displayed here.


+ Add Widget

Select Widget

Network Availability
Alarms
Locator
Favorites
Quarantine Manager
Licenses
Top N Clients Detail View
Top N Switches Summary View
Top N Ports Utilization Summary View
Top N Ports Utilization Detail View
Top N Ports Utilization Trending View
Top N Apps Summary View
Top N Apps Detail View

Top N Clients Detail View

UNPs statistic by bytes count



UNPs


Example Only

More


Keep track of the number of flows seen for each user and report on the top N Clients.

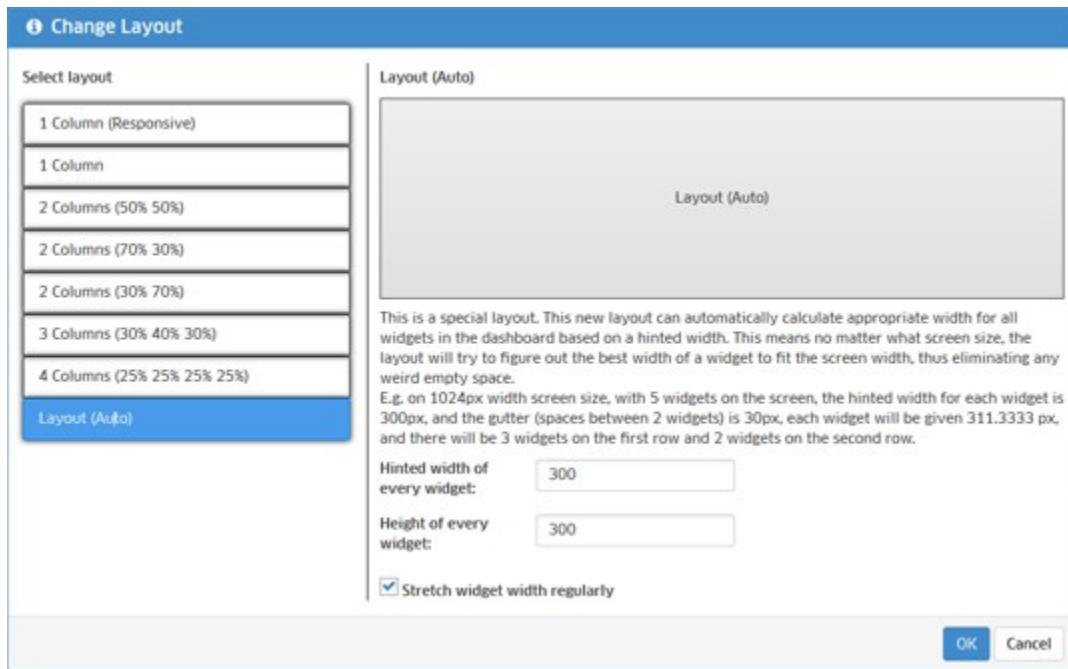
When you select a widget, a sample and summary of widget information is displayed.

Removing Widgets

To remove a widget from the dashboard, click on the Delete icon (x) at the top-right corner of the widget. You can add the widget back or add a new widget by clicking on the Settings icon  and selecting **Add Widget**.

Changing the Dashboard Layout

The Settings option also enables you to customize the Dashboard layout. To customize the layout, click on Settings icon  and select **Change Layout**. The Change Layout Screen appears. By default, widgets are displayed in the "Auto" layout. To change the layout, select an option on the left. When you select a layout on the left, the layout is displayed in the Layout area. To choose that layout, click **OK**. The change takes effect immediately. If necessary, you can re-arrange the widgets in the new layout, by clicking and dragging them to new locations on the Dashboard.



Application Widgets

The following application widgets are available. Some widgets (e.g., Network Availability, Licenses) collect and display information automatically. Others (e.g., Alarms, Top N Clients) require configuration in an application (e.g., Notifications, Application Visibility/Analytics) to display information. The descriptions below provide an overview of the information provided by each widget and a link to the applicable help file for application configuration, if required. For detailed configuration procedures, see the applicable online help.

Alarms

Displays an overview of Notification alarms for all discovered network devices by severity level (Normal, Warning, Minor, Major Critical). Click on a severity level to view information broken down by device. You must configure traps in the Notifications application to generate and display information for this widget.

Application Count Detail View

Uses Signature Files to display detailed information about the top applications being accessed on the network for OS6860/OS6860E Switches on byte/packet count. The information is displayed in a line chart view. You must create a Signature Profile in the Application Visibility application and assign the profile to switches/ports to generate and display information for this widget.

Application Count Summary View

Uses Signature Files to display summary information about the top applications being accessed on the network for OS6860/OS6860E Switches based on byte/packet count. The information is displayed in a pie chart view. Hover over a section to view application information. Click on a section to view the source switch. You must create a Signature Profile in the Application Visibility application and assign the profile to switches/ports to generate and display information for this widget.

Application Count - UNPs Detail View

Uses Signature Files to display detailed information about the top applications being accessed on the network for OS6860/OS6860E Switches based on byte/packet count. The information is filtered by UNP Profile (UNP Access Role Profile) information and is displayed in a bar chart view.

Application Count - UNPs Summary View

Uses Signature Files to display summary information about the top applications being accessed on the network for OS6860/OS6860E Switches based on byte/packet count. The information is filtered by UNP Profile (UNP Access Role Profile) information and is displayed in a pie chart view.

Application Discovery OS6860

Uses Signature Files to display information about the top applications being accessed on the network based on Signature Profiles for OS6860/OS6860E Switches based on flow information (number of flows). Hover over a section to view application information. Click on a section to view the source switch. You must create a Signature Profile in the Application Visibility application and assign the profile to switches/ports to generate and display information for this widget.

Application Discovery OS6900/OS10K

Uses Signature Files to display information about the top applications being accessed on the network for OS10K/OS6900 Switches based on byte/packet count. Hover over a section to view application information. Click on a section to view the source switch. You must create a Signature Profile in the Application Visibility application and assign the profile to switches/ports to generate and display information for this widget.

Favorites

The Favorites widget allows users to create a list of “quick links” for easy access to any OmniVista 2500 NMS application. To add an application to the Favorites list, start typing the name of the application in the text box. As text is entered, auto-complete displays a list of matching applications. Click an application to select it, then click the add icon (+). To replace an application in the list of Favorites, click the **Edit** icon. Begin typing the new application name. Select the application and click the **Save** icon. Once an application has been added to the Favorites widget, a link to the application will also display under the Favorites tab in the OmniVista 2500 NMS Main Navigation Menu. To remove an application from the Favorites widget, click the Delete icon next to the application link.

Licenses

Displays an overview of License usage. Information is displayed for all enabled licenses (e.g., AOS, VM), including the number of devices being managed and expiration (if applicable). Click on a license type to view the number of devices being managed and maximum number of devices that can be managed under the license. This widget automatically displays information for all discovered network devices.

Locator

Used to locate the switch and slot/port that is directly connected to a user-specified end station. Enter the end station's IP Address, Host Name, MAC address, or Authenticated User ID to locate the switch and slot/port to which the end station is connected. Enter the search criteria and click on **Locate**. OmniVista will display the results in the Locator application.

Network Availability

Displays the current operational state of all discovered network devices (Up/Warning/Down). Each category is displayed as a percentage of all monitored switches. This widget automatically displays information for all discovered network devices. Click on a warning state to view a list of devices in that state. You can then click on a link to the Notifications application to view traps. This widget automatically displays information for all discovered network devices.

ProActive Lifecycle Management

Provides an overview of ProActive Lifecycle Management information for network devices, such as the percentage/number of network devices under warranty or out of warranty. Information includes: OS Release Support, Hardware Support, Warranty Support, and Service Support. Page through the widget to view each category. Click on the widget to go to the ProActive Lifecycle Management website. This widget is only available if your OmniVista ID is registered with the ALE ProActive Lifecycle Management Portal.

Quarantine Manager

Displays an overview of the number of devices in each Quarantine Manager category (Candidates, Banned, Never Banned, Disabled Ports). You must configure Quarantine Manager to display information for this widget.

Top N Apps Detail View

Displays a detailed bar chart view of the top applications being accessed on the network for the configured time period. For example, if a report is configured to show data for the last 24 hours, the Detail View displays data for each hour within those 24 hours. Hover over an area in a bar chart to view information about an application. You must create an Analytics Profile in the Analytics application and assign the profile to switches/ports to generate and display information for this widget. The Top N Applications are determined using sFlow. sFlow is enabled on a port when you create an Analytics Profile. sFlow can also be enabled on a port using the CLI. If sFlow is configured on a port using the CLI and the OmniVista Server is configured as the receiver, Top Applications data will be displayed in OmniVista.

Top N Apps Summary View

Displays summary information about the top applications being accessed on the network in a pie chart view with each application displayed as a percentage of the total traffic for all monitored switches. Hover over a section to view application information. Click on a section to view client or switch information. You must create an Analytics Profile in the Analytics application and assign the profile to switches/ports to generate and display information for this widget.

Top N Clients Detail View

Displays detailed bar chart view of the top network clients. For example, if a report is configured to display data for the last 24 hours, the Detail View displays data for each hour within those 24 hours. Hover over an area in a bar chart to view information about a client. You must create an Analytics Profile in the Analytics application and assign the profile to switches/ports to generate and display information for this widget. The Top N Clients are determined using sFlow. sFlow is enabled on a port when you create an Analytics Profile. sFlow can also be enabled on a port using the CLI. If sFlow is configured on a port using the CLI and the OmniVista Server is configured as the receiver, Top Client data will be displayed in OmniVista.

Top N Clients Summary View - Displays summary information for the top network clients based on the number of traffic flows for each client in a pie chart view. Hover over a section to view client information. Click on a section to view application information for that client. You must create an Analytics Profile in the Analytics application and assign the profile to switches/ports to generate and display information for this widget.

Top N Ports Utilization Detail View

Displays a detailed bar chart view of the top network ports based on utilization. While the Summary View displays the information for the configured time period (e.g., last 24 hours), this view provides a detailed view of the specified time interval. For example, if a report is configured to display information for the last 24 hours, the Detail View displays information for each hour within those 24 hours. Hover over an area in a bar chart to view information about devices/ports. You must create an Analytics Profile in the Analytics application and assign the profile to switches/ports to generate and display information for this widget.

Top N Ports Utilization Summary View

Displays a summary view of the top network ports based on utilization. Switches/ports are displayed in a list view from highest to lowest utilization for the configured time period (e.g., day, week). You must create an Analytics Profile in the Analytics application and assign the profile to switches/ports to generate and display information for this widget.

Top N Ports Utilization Trending View

Displays predicted future port utilization in a line chart based on past utilization. Port utilization predictions can provide valuable insight for capacity management. Click on a data point to view information about devices/ports. You must create an Analytics Profile in the Analytics application and assign the profile to switches/ports to generate and display information for this widget.

Top N Switches Detail View -

Displays detailed bar chart view for the top switches on the network in terms of the switch's resource usage. Resource usage is calculated based on switch's CPU usage, memory usage, and temperature. The switches are displayed in order on each bar chart from highest to lowest utilization. This widget automatically displays information for all discovered network devices.

Top N Switches Summary View

Displays a list view of the top switches on the network in terms of the switch's resource usage from highest to lowest utilization. Resource usage is calculated based on switch's CPU usage, memory usage, and temperature. This widget automatically displays information for all discovered network devices.

3.0 Analytics

The Analytics Application provides users with a comprehensive view of network resource utilization, including views of users, devices, and applications. The application also provides information on usage trends, including predictive analysis of future network resource utilization. The Reports Screen (shown below) is used to view Analytics Reports and configure how the information is displayed.



Note: The Analytics application provides real-time viewing of Analytics Reports. You can also schedule Analytics Reports to be generated and stored as PDF documents using the Report application. This way, in addition to real-time viewing in the Analytics application, you can automatically generate and store Analytics Reports that you can view at any time.

Using Analytics

The Analytics application enables users to create different reports (e.g., Top N Applications, Top N Ports Utilization) that provide a comprehensive view of network and device utilization. The following screens are used to view/analyze the network using the Analytics application:

- **Reports** - Used to configure reports that provide a comprehensive view of network resource utilization and device status. Top N Applications, Top N Applications - Advanced, Top N Clients, Top N Switches, and Top N Ports Utilization Reports can be configured to show network utilization over different time periods (e.g., daily, hourly, monthly), and show trends in network utilization over those time periods.

The Top N Ports Utilization Report can also provide predictive analytics to show expected future usage. Other reports can provide a "real-time" view of all discovered network switches (Network Availability, Alarms). The following reports can be created:

- **Top N Applications** - Displays information about the top applications being accessed on the network, including which users are using an application, and which switches have the most traffic for an application. The Top N Applications are determined using sFlow. To generate a Top N Applications Report, you must first create an Analytics Profile.
- **Top N Applications - Advanced** - Displays information about the top applications being accessed on the network based on Signature Profiles configured in the Application Visibility application.
- **Top N Clients** - Displays information for the Top Network clients including the number of traffic flows for each client. To generate a Top N Clients Report, you must first create an Analytics Profile.
- **Top N Switches** - Displays information for the top devices on the network in terms of the device's resource usage. Devices are ranked based on the device's CPU usage, memory usage, and temperature.
- **Top N Ports Utilization** - Displays network ports by utilization over time; and also provides predictive analytics to show future port utilization trends. To generate a Top N Ports Utilization Report, you must first create an Analytics Profile.
- **Network Availability** - Displays the current operational state of all discovered network devices (Up/Warning/Down).
- **Alarms** - Displays network alarms by severity level for all discovered network devices.
- **SIP Active Calls** - Displays Active Call Record data for selected SIP-enabled switches. This report is generated by the SIP application and is displayed in table format only. This report
- **SIP Ended Calls** - Displays Ended Call Record data for selected SIP-enabled switches. This report is generated by the SIP application and is displayed in table format only.
- **Performance Monitoring** - Used to collect, monitor, and view statistical information for devices throughout the network.
- **Profiles** - Used to create Analytics Profiles. To generate a Top N Applications, Top N Clients, and Top N Ports Utilization Report you must first create an Analytics Profile that defines the switches/ports that you want to view and the type of information that you want to view on those switches/ports.
- **Summary View** - Displays basic information on all supported network devices, including any Analytics Profiles defined for a device.
- **Applications Management** - When generating a Top N Applications Report, the Analytics application uses port numbers to identify application traffic. This screen is used to create port/application mappings to identify applications traffic.
- **Anomalies** - Displays any port utilization anomalies. An anomaly is a utilization data point that fall outside of expected norms based on past usage.
- **Settings** - Used to configure preferences for port utilization trending and anomaly detection in the Analytics application.

Note: Remember, to generate Top N Applications, Top N Clients, and Top N Ports Utilization Reports, you must first create an Analytics Profile that defines the switches/ports that you want to view and the type of information that you want to view on those switches/ports. To generate a Top N Applications - Advanced Report, you must first create a Signature Profile in the Application Visibility application that defines the switches/ports that you want to view and the type of information that you want to view on those switches/ports. Data will only be gathered and displayed for those switches/ports included in the profile. You do not need to create a profile for Network Availability, Alarms, Top N Switches, or SIP Reports.

Configuring Analytics

The first step in generating analytics information for Top N Applications, Top N Clients, and Top N Ports Utilization Reports is to go to the Profiles Screen and create an Analytics Profile. Analytics information is gathered by creating an Analytics Profile that specifies the information to be viewed (e.g., Top N Applications, Top N Ports Utilization) and the network switches/ports that will be monitored. The Profile Type will determine the type of Analytics Report that you can generate (e.g., Top N Applications, Top N Users). Reports will generate data only for those switches/ports included in a profile.

Network Availability, Alarms, Top N Switches, and SIP Reports provide a "real time view" of the network. You do not need to create a profile for these reports. However, to view network alarms (Alarm Report) you must go to the Notifications application and configure traps on the switches you want to monitor. Network alarms will then be displayed on the Alarms Report Screen. (These alarms are also displayed, along with all network alarms, in the Notifications application.)

You can view all Analytics Reports on the applicable report screen. The information in the reports is presented in graphical and linear format, depending on the report type.

Reports

Analytics Reports provide users with a comprehensive view of network resource utilization, including information on users, devices, and applications. Reports can also provide information on usage trends, including predictive analysis of future network resource utilization. Top N Applications, Top N Applications - Advanced, Top N Clients, Top N Switches, and Top N Ports Utilization Reports can be configured to show network utilization over different time periods (e.g., daily, hourly, monthly), and show trends in network utilization over those time periods. The Top N Ports Utilization Report can also provide predictive analytics to show expected future usage. Other reports can provide a "real-time" view of all discovered network switches (Network Availability, Alarms). You can view the reports in different formats and customize how the data is displayed. The following reports can be created:

- **Top N Applications** - Displays information about the top applications being accessed on the network, including which users are using an application, and which switches have the most traffic for an application.
- **Top N Applications - Advanced** - Displays information about the top applications being accessed on the network based on Signature Profiles configured in the Application Visibility application.
- **Top N Clients** - Displays information for the Top Network clients including the number of traffic flows for each client.
- **Top N Switches** - Displays information for the top devices on the network in terms of the device's resource usage. Devices are ranked based on the device's CPU usage, memory usage, and temperature.
- **Top N Ports Utilization** - Displays network ports by utilization over time; and also provides predictive analytics to show future port utilization trends.
- **Network Availability** - Displays the current operational state of all discovered network devices (Up/Warning/Down).
- **Alarms** - Displays network alarms by severity level for all discovered network devices.
- **SIP Active Calls** - Displays Active Call Record data for selected SIP-enabled switches. This widget provides a link to the SIP Active Calls Screen in the SIP application. The reports are generated by the SIP application and are displayed in table format only.
- **SIP Ended Calls** - Displays Ended Call Record data for selected SIP-enabled switches. This widget provides a link to the SIP Ended Calls Screen in the SIP application. The reports are generated by the SIP application and are displayed in table format only.

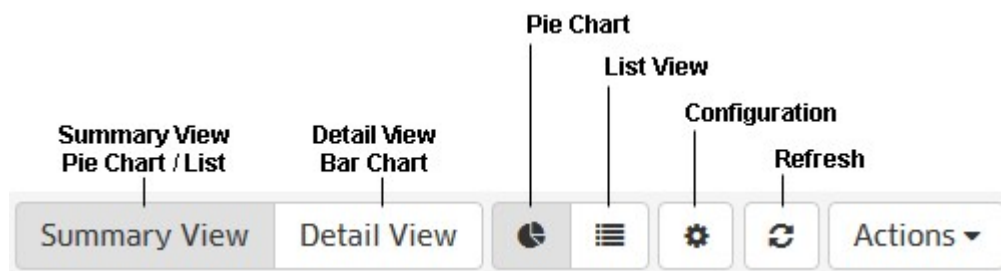
Note: To generate a Top N Applications, Top N Applications - Advanced, Top N Clients, and Top N Ports Utilization Report, you must **first create an Analytics Profile** using the Profile Screen that defines the switches/ports that you want to view and the type of information that you want to view on those switches/ports. Data will only be gathered and displayed for those switches/ports included in the profile. You do not need to create a profile for Network Availability, Alarms, Top N Switches, or SIP Reports. These reports simply show real-time information for all discovered switches.

Note: Top N Apps & Clients Profiles use sFlow to gather information. When these profiles are created, the OmniVista Server is automatically configured as the sFlow Receiver. However, sFlow can be configured on a device outside of OmniVista (e.g., using the CLI). If sFlow is configured on a device outside of OmniVista and the OmniVista Server is designated as the sFlow Receiver, the information for that device is sent to OmniVista and included in Top N Applications and Top N Clients Reports. (Information will be displayed in these reports even if no profile was created and assigned in OmniVista.) If the device is not known to OmniVista (or if the Analytics application is not supported on the device), sFlow information is sent to OmniVista, but the information is not included in those reports.

The sections below describe the different report options and basic behavior for all reports. The report options vary depending on the report type (e.g., Top N Applications, Top N Switches). Specific views/options are detailed in the help pages for each report type. Click on a link above to view specific instructions for each report type.

Report Options

Analytics Reports can be viewed in different formats (e.g., pie chart, bar chart). You can also configure a custom view to change the amount of information displayed (e.g., the number of Applications/Clients displayed), as well as the timeframe that you want to view (e.g., last 24 hours, last 7 days). You can also view data trends by "drilling down" in a Detailed Report. You can also print a report or download a report in PDF or PNG format, and even include the data as part of a scheduled report that is automatically generated in the Report application. These options can be configured using the Options Bar (shown below) displayed at the top of every report.



View Options

By default, the Summary View is initially displayed for all reports. This may be displayed graphically as a pie chart (e.g., Top N Applications, Top N Ports Utilization) or in a list (e.g., Top N Switches). The Detail View displays a detailed subset of the information in a bar chart format. While in the Detail View, you can also display an even more detailed subset of the data to view data trends. For example, if a Summary View is displaying data for the last week, the Detail view will display data for each day of the last week; and clicking on a day in the Detail view chart will display data for each hour of that day, enabling you to view hourly data trends.


For "Top N" Reports (e.g., Top N Applications, Top N Clients), you can configure the amount of data displayed and the time period you want to view. You can set the number of "top" (in terms of utilization) applications, clients, or switches you want in the display. For example, you might want to see the top 10 applications displayed in a Top N Applications Report; or the top 20 ports displayed in a Top N Ports Utilization Report. You can also configure the time period to display (last 24 hours, last 7 days, last 4 weeks).

The chart legend to the right of each display labels each item in the chart by color and text. By default, information from all switches/ports included in a profile is displayed. The number of entities displayed in the legend and the chart depends of the number you configure for the profile (e.g., top 10, top 20). However, you can click on the Select Devices button at the top of the screen to display only information from specific switches/ports.

Note: You may notice a category labeled "Others" in "Top N" Reports. Remember, only the "top" applications, clients, or switches as determined by the profile (top 10, top 20) are displayed. There may be many others in the profile that are not in the "top" 10 or 20. The "Others" category gives you an idea of all of the other applications, clients, or switches in the profile with low utilization rates that do not qualify as a "top" application, client, or switch.

Summary View

The Summary View displays information either in pie chart format, with each entity (e.g., application, client) displayed at a percentage of the total for the configured time period (e.g., 24 hours); or in list format, with each entity listed from highest to lowest. By default, data for the past 24 hours is shown. However, you can change the timeframe, as well as the number of entities displayed (top applications, top ports). The examples below show an Hourly Summary View for the Top N Applications Report.



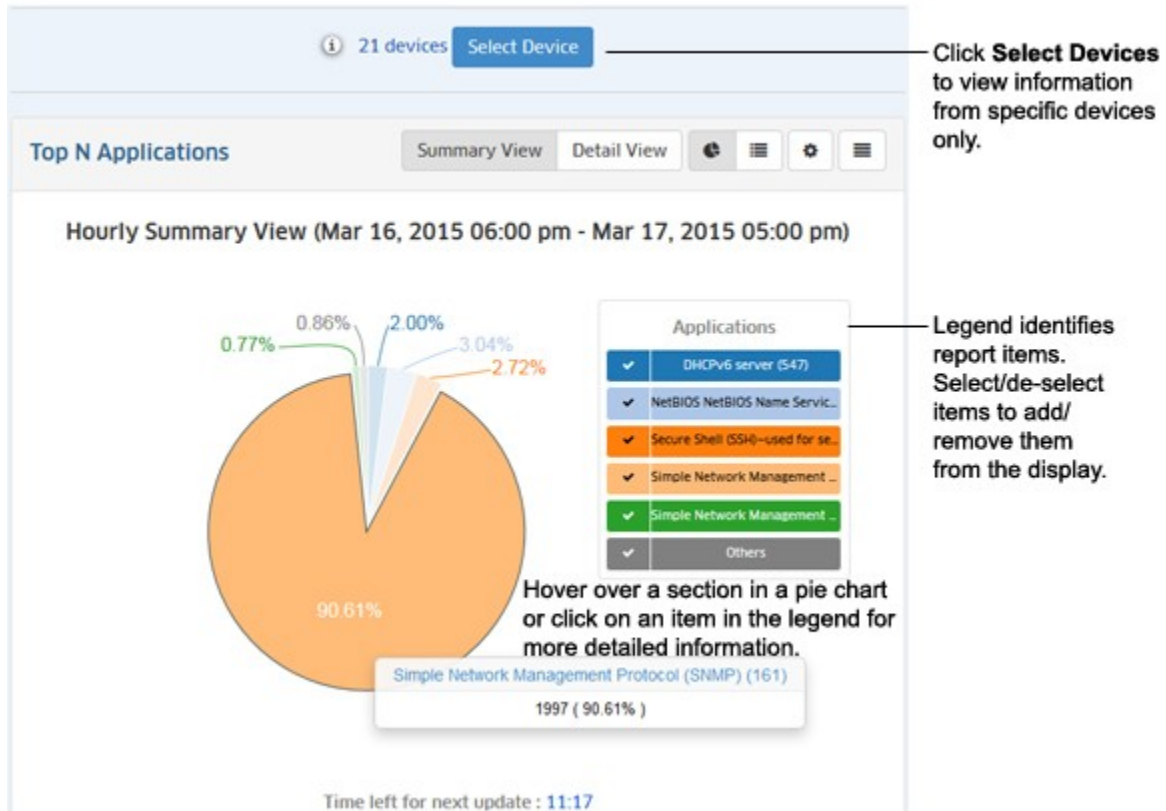
Summary View List



Pie Chart Format

The Pie Chart Format displays information in a pie chart with each entity displayed as a percentage of the total. The legend identifies each item in the chart by color and text. For example, the legend in a Top N Applications Report (shown below) identifies the applications displayed in the pie chart. (The legend in a Top N Ports Utilization Report would identify the switches/ports displayed.) You can hover over a section of the chart (or click on an item in the legend) to view detailed information for that section. For example, in the Top N Applications pie chart below, you could hover over an application in the chart to view the number of flows from that application. You can also click select/deselect an item in the legend to add/remove the item from the display.

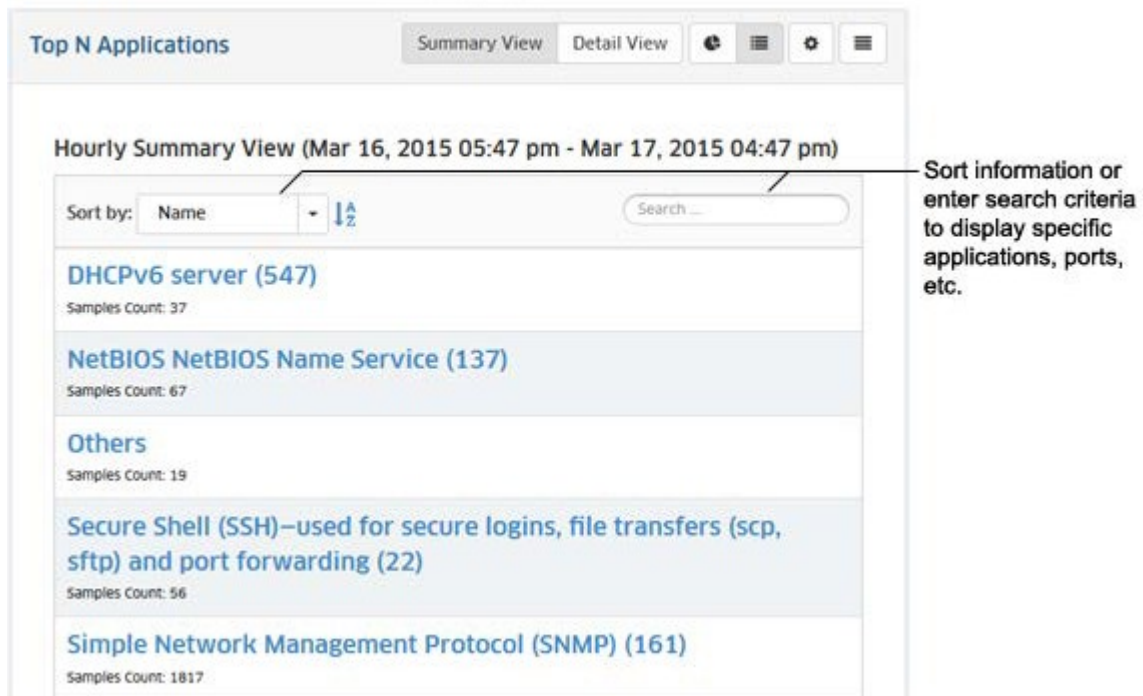
Pie Chart Format



List Format

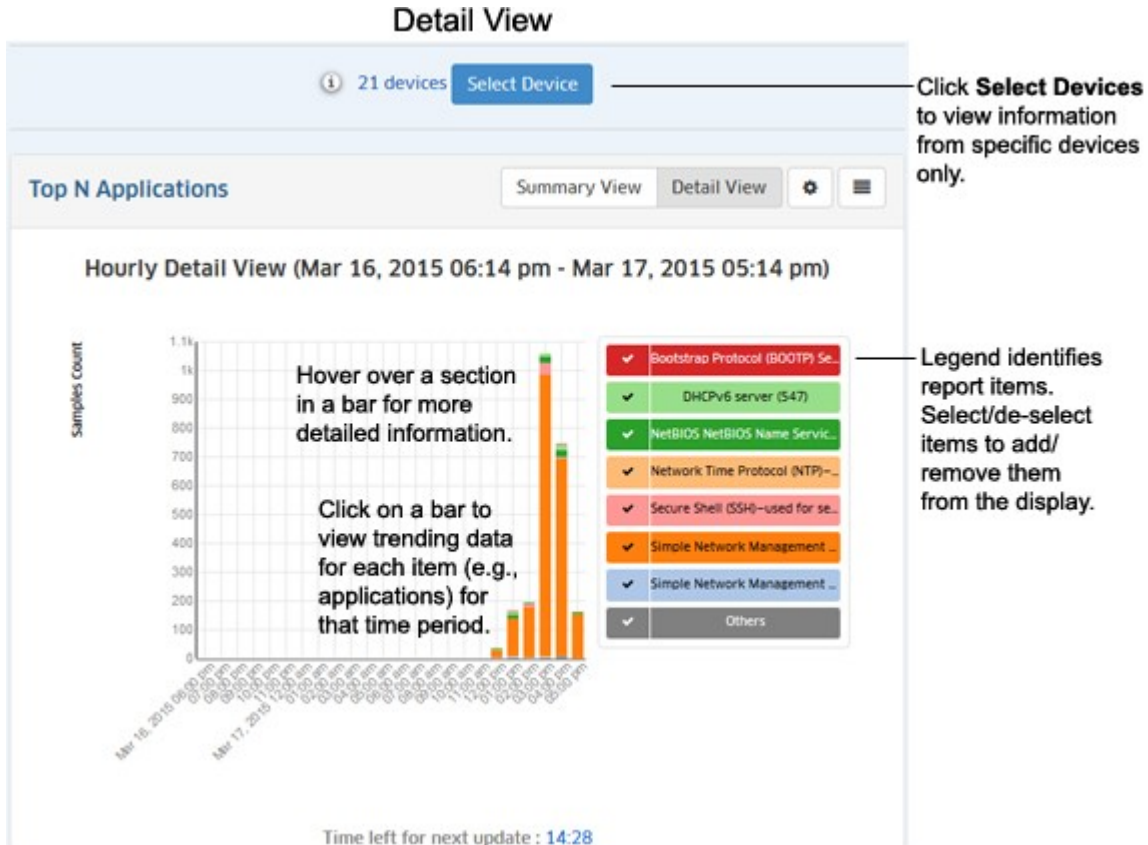
The List Format displays information in list form (e.g., a list of Switches, Applications, or Users displayed from highest to lowest). You can sort the display by specific criteria (e.g., name), sort in ascending/descending order by, or search for and display specific information (e.g., display only a specific application or port number). The example below shows a Top N Applications Report.

List Format



Detail View


The Detail View displays a detailed subset of the information in bar chart format for the configured time period. For example, if a report is configured to show data for the last 24 hours, the bar chart view would display data for each hour over those 24 hours. The legend identifies each item in the chart by color and text. For example, the legend in a Top N Applications Report (shown below) identifies the applications displayed in the bar chart. (The legend in a Top N Ports Utilization Report would identify the switches/ports displayed.) You can hover over an area in a bar to view detailed information for that item. For example, in the Top N Applications bar chart below, you could hover over an area in a bar to view the number of flows from that application. Or you can click on an item in the legend to isolate the item in the display and show the same detailed information. You can also click select/deselect an item in the legend to add/remove the items from the display. You can also view data trends by "drilling down" on a data set to see a subset of that data.



Data Trends

You can view data trends by "drilling down" on a data set to see a subset of that data. Click on a bar in the chart to view the data trend for that selection. For example, if you selected one of the bars in an Hourly Detail View, the trend for that hour would be displayed in 15 minute increments (as shown below). (If you selected one of the bars in a Daily Detail View, the trend for that hour would be displayed in one-hour increments.) You can scroll forward or back through the trending date using the arrows at the top of the chart.


Trend View



Depending on the timeframe in the chart (e.g., Monthly, Weekly), data trend subsets are displayed as follows:

- Monthly Details View - A Weekly Trending View
- Weekly Details View - Daily Trending View Daily
- Details View - Hourly Trending View
- Hourly Details View - 15-minutes Trending View.

Configure a Custom Report

You can configure the report view by clicking on the Configuration icon  and customizing the display. The Configuration Screen for the Top N Applications Report is shown below. The configuration options vary depending on the report type. Specific field descriptions are defined in the help pages for each report type. Once you update any options and click on the **Save** button, reports will be displayed in the new format.

Configuration Screen

The screenshot shows a configuration window with the following fields and controls:

- Default Devices:** 0 device(s) selected, with a "Select Device" button.
- Number of Top Applications:** A text input field containing "5" and up/down arrow buttons.
- Interval Type:** A radio button interface with "Up Until Now" selected and "Custom" as an alternative.
- Time Interval:** A text input field containing "24 hours" and a dropdown arrow.
- Auto Refresh Timer:** A text input field containing "15", a "min(s)" label, and up/down arrow buttons.
- Buttons:** "Save" and "Cancel" buttons at the bottom right.

Download/Print a Report

You can download a report in PDF or PNG format or send the report to a printer by clicking on the **Actions** button in the Options Bar and making a selection from the drop-down menu (**Download PNG Image/Download PDF Document/Print Image**).

Schedule a Report

You can add the current report view to a Report that you create in the Report application by clicking on the **Actions** button in the Options Bar and selecting **Add to Report**. The Report application enables you to create and schedule Analytics Reports that can be viewed and stored as PDF documents. This way, in addition to real-time viewing of Analytics Reports in the Analytics application, you can automatically generate and store Analytics Reports that you can view at any time. See the Report configuration Help for more information.

Top N Applications

The Analytics Top N Application Report Screen displays information about the top applications being accessed on the network. The Top N Applications are determined using sFlow. OmniVista identifies the applications using the TCP/UDP port obtained from sFlow packets. In other words, traffic on a specific port is identified as coming from a specific application. Well known ports (e.g., 161 for SNMP, 80 for HTTP) are automatically identified and labeled in the Top N Applications Report. Other applications can be mapped to a port using the Applications Management Screen.

To generate a Top N Applications Report, you must first create and assign an Analytics Profile using the Profile Screen that defines the switches/ports that you want to view and the type of information that you want to view on those switches/ports.

Note: sFlow is enabled on a port when you create an Analytics Profile. However, sFlow can also be enabled on a port using the CLI. If sFlow is enabled on a port using the CLI and the OmniVista Server is configured as the receiver, Top N Applications data will be displayed in OmniVista.

Note: sFlow packets cannot be sent through the EMP Port. If you want to gather Top N Application data from a switch you cannot use the EMP IP when discovering the switch.

By default, the Summary View is displayed (pie chart) with each application displayed as a percentage of the total number of flows for the configured time interval (e.g., last 24 hours). Information from all switches in the profile is displayed. However, you can click on the **Select Devices** button to display only information from specific switches. The information can be displayed in different formats, and you can also configure the amount of information displayed.

Report Views

The Top N Applications Report can be displayed in a Summary View or a Detail View. The Summary View provides a summary of application traffic for the configured time interval (e.g., last 24 hours (default), last 7 days). The Detail View displays a subset of the data in a bar chart format. For example, if a report is configured to display data for the last 24 hours, the Summary View will display a summary of the data for the last 24 hours; and the Detail View will then display data for each hour within those 24 hours.

Note: Report Views and configuration options are configured using the Options Bar located at the top of the report. This help page contains view and configuration information specific to Top N Applications Reports. For specific information on all of the options available, see the "Report Options" section of the Analytics Reports Help.

Summary View

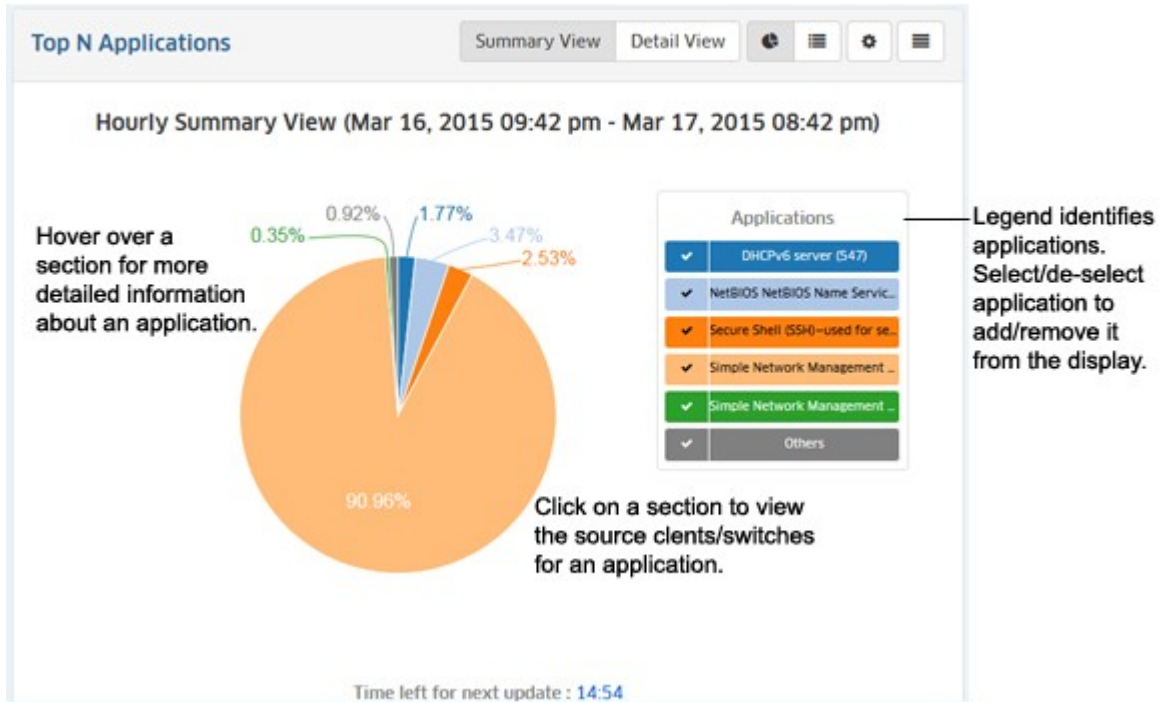
By default, the Summary View is displayed. This view provides a summary of application traffic for the configured time interval (e.g., last 24 hours). By default, the pie chart format is displayed; however a list format is also available.

Pie Chart Format

By default, the Summary View is displayed as a pie chart, with each application displayed as a percentage of the total traffic for all monitored switches for the configured time interval. By default, an Hourly Summary Report is displayed, showing a summary of the data over the last 24 hours. (The Detail View will then display detailed information for each hour.) However, you can configure the report to display different time intervals (e.g., last 24 hours, last 7 days).

The legend on the right of the screen identifies each application in the chart by color and text. You can hover over a section of the chart to view detailed information for that section. Or you can click on an application in the legend to isolate the item in the display and show the same detailed information. You can also click select/deselect an application(s) in the legend to add/remove the application(s) from the display. The example below shows an Hourly Summary View.

Top N Applications Summary View - Pie Chart



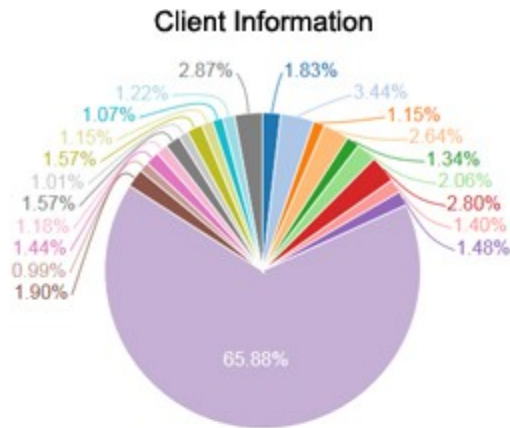
Hover the mouse over a section of the chart (or click on an application in the legend) to view the number of flows for that application over the time interval displayed (e.g., last 24 hours). In the example below, hovering over the SNMP section of the pie chart shows the total number of SNMP flows as 4851, or 90.34% of the totals number of application flows. You can also view information on which clients are accessing an application, and which switches are passing traffic for that application.




Client and Switch Information

When in the Pie Chart View of the Top N Applications Report you can view information about clients accessing an application (by source IP address) or the switches passing the application traffic. Click on a section and select Clients to view client information, or Switches to view switch information. The pie chart will be broken down by client or switch for that application (as shown below). The legend identifies the client or switch by color and text, or you can hover over a section to view the client/switch IP address (along with detailed flow information). You can also select/de-select switches/clients in the legend to add/remove them from the display.

The example below shows client and switch information for the SNMP application. In this example, many clients are using the SNMP application with the traffic passing between two switches. Click on the Back Arrow (<) above the legend to return the default view.



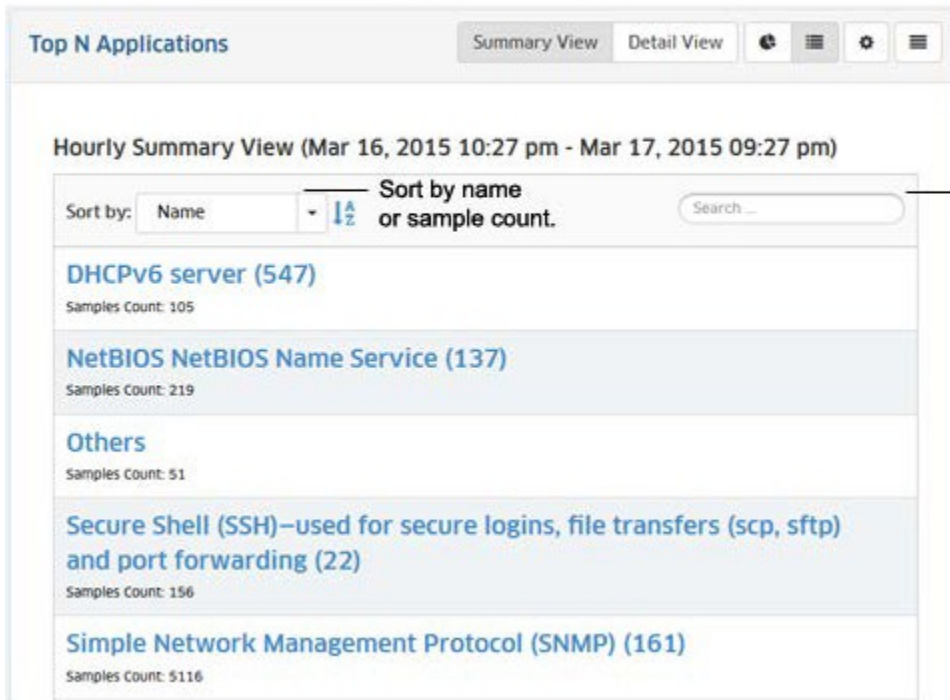
Information for each client/switch is displayed.
Hover over a section for more detailed information.



List Format

The list format displays a list of applications with packet count information for each one. By default, the list is displayed by application name in alphabetical order; however you can select "Samples Count" in the **Sort by** drop-down menu to display the applications by sample count. You can also search for and display a specific application by entering the application name in the **Search** field.

Top N Applications Summary View - List



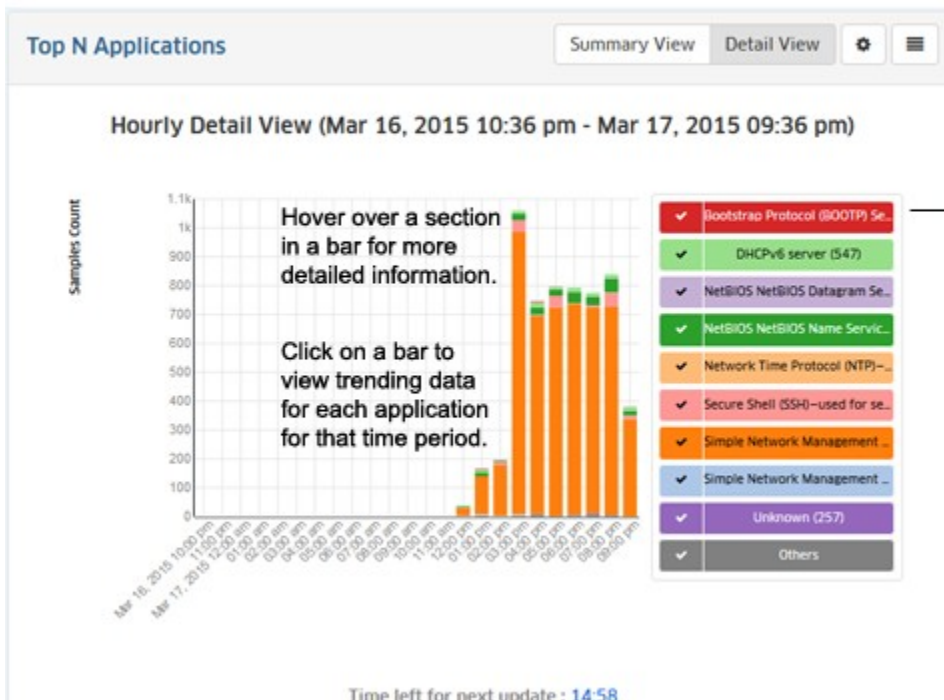
Search for a specific application.

Detail View

The Detail View displays a detailed subset of the information in bar chart format for the configured time period. For example, if a report is configured to show data for the last 24 hours, the Summary View will display a summary of the data for the last 24 hours; and the Detail View will then display data for each hour within those 24 hours.

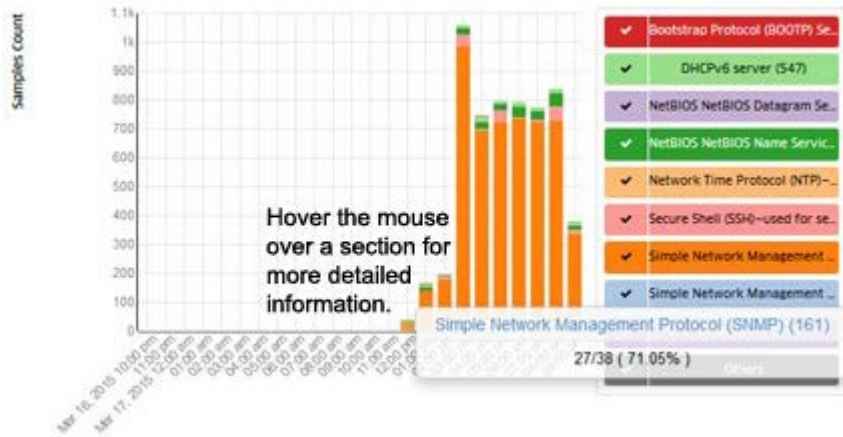
Note: You can also click on a bar to view usage trends for that time interval. For example, if you clicked on a day in the chart below, you can view hourly usage trends for each application for that day.

Top N Applications Detail View




Legend identifies applications. Select/de-select applications to add/remove them from the display.

As in the Summary Pie Chart view, you can hover the mouse over a section of the chart to view the number of flows for that application over the configured time interval (e.g., hour). In the example below, hovering over the SNMP section of a bar shows the total number of SNMP flows as 27 out of a total of 38 flows for that hour, or 71.05% of the total number of application flows for that hour.




Trending Information

When in Detail View, you can click on a bar in the chart to view usage trends for each application for the selected time interval by "drilling down" on a data set to see a subset of that data. For example, if you selected one of the bars in an Hourly Detail View, the trend for that hour would be displayed in 15 minute increments (as shown below). Click on a data point in the trending view for more detailed information. You can scroll forward or back through the trending date using the arrows at the top of the chart.



Configuring the Information Displayed

You can configure the amount of information displayed (e.g., the number of applications you want to view) as well as the time interval that you want to view. To configure the report display, click on the Configuration icon  to bring up the Configuration Screen, then complete the fields as described below to configure how information displayed in the report.

- **Default Devices** - By default, all top switches/ports are displayed. However, you can click on the **Select Devices** button to display only information from specific switches.
- **Number of Top Applications** - The number of applications you want to display (Range = 1 - 20, Default = 5).
- **Interval Type** - The time interval for the information:
 - **Up Until Now** - Displays all information in the selected time interval (e.g., last 24 Hours).
 - **Custom** - Set the start and end time for the information you want to display. You can display up to 3 months of data. When data reaches the 3-month maximum, it is overwritten with new data.
- **Time Interval** - The time interval you want to display in the report (the past 24 Hours, 7 Days, or 4 Weeks).
- **Auto Refresh Timer** - How often you want to refresh the data display, in minutes (Range = 15 - 60, Default = 15). The configuration option is only available when "Up Until Now" is selected for Interval Type.

When you are done, click the **Save** button. The report display will immediately change to the new view. This will remain the view until it is changed again.

Top N Applications - Advanced

The Analytics Top N Applications - Advanced Report Screen displays information about the top applications being accessed on the network based on Signature Profiles configured in the Application Visibility Application. Signature Profiles include the specific application groups/applications being monitored as well as the specific switches being monitored, so the information displayed is determined by the applications and switches included in the profile. Only information for those applications and switches is displayed. By default, the Summary View is displayed (pie chart) with each application displayed as a percentage of the total for the configured time interval (e.g., last 24 hours). The information can be displayed in different formats, and you can also configure the type and amount of information displayed.

The Application Visibility Application and Top N Applications - Advanced Reports are supported on OS10K Switches (AOS 7.3.4.R01 and later), OS6900 Switches (AOS 7.3.4.R02 and later) and OS6860/6860E Switches (AOS 8.2.1.R01 and later).

Note: To generate a Top N Applications - Advanced Report, you must first create a Signature Profile in the Application Visibility application and assign the profile to switches/ports that you want to view.

Note: Report Views and configuration options are configured using the Options Bar located at the top of the report. This help page contains view and configuration information specific to Top N Applications - Advanced Reports. For specific information on all of the options available, see the "Report Options" section of the Analytics Reports Help.

Report Views

Signature Profiles are created in the Application Visibility Application based on supported switch type (OS10K/OS6900 or OS6860/OS6860E). When you click on the Top N Applications - Advanced link, you have the option to select the type of information you want to display (App Discovery or App Count). You can also select to view information for all switches in a type, or select specific switches.

- App Discovery - Displays traffic flow information for applications/application groups discovered on the network, and the percentage of network resources being used by each application for the selected switches and configured time period. OS10K/OS6900 Switches provide byte/packet information; not flow information. OS6860/OS6860E Switches provide flow information (number of flows) in the App Discovery view and packet/byte information in the App Count view. All switch types sample data.
 - *For all 6900s/OS10Ks* - Displays flow information for all OS6900/OS10K Switches.
 - *For all 6860s* - Displays flow information for all OS6860/6860E Switches.
 - *Manually select devices* - Click on the link, then click on the **Select Device** button to bring up the Device Selection Window. Select a switch type from the Select OS drop-down menu (OS6860, OS6900/OS10K), select the switch(es) you want to include in the report, and click **OK**. You can only include switches from the same device type in the report. In other words, you can only select either OS6860 or OS6900/OS10K switches.
- App Count (Supported on OS6860/6860E Only) - Displays packet/byte count information for applications/application groups discovered on the network over the configured period of time, and the percentage of network resources being used by each application for the selected switches and configured time period. You can also view application information by UNP Profile.
 - *Manually select devices* - Click on the link, then click on the **Select Device** button to bring up the Device Selection Window. Select a switch type from the Select OS drop-down menu

(OS6860, OS6900/OS10K), select the switch(es) you want to include in the report, and click **OK**. You can only include switches from the same device type in the report. In other words, you can only select either OS6860 **or** OS6900/OS10K switches.

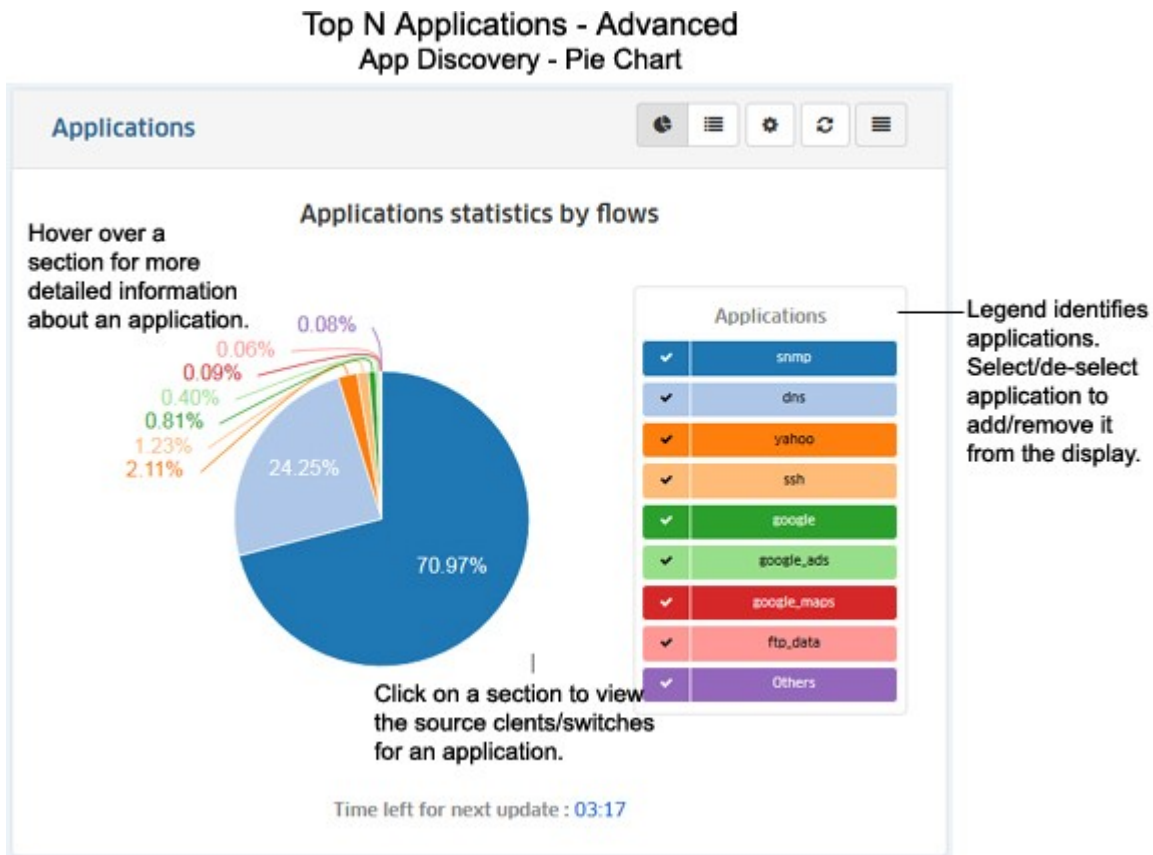
- o *For all Devices* - Displays packet/byte count information for all switches.

App Discovery


The App Discovery view displays traffic flow information for applications/application groups discovered on the network, and the percentage of network resources being used by each application for the selected switches and configured time period. The information can be displayed in a Pie Chart Format (Default) or in a List Format.

Pie Chart Format

By default, App Discovery reports are displayed as a pie chart, with each application displayed as a percentage of the total traffic for the selected switches. By default, an Hourly Summary Report is displayed, showing a summary of the data over the last 24 hours. However, you can configure the report to display different time intervals (e.g., last 24 hours, last 7 days). The legend on the right of the screen identifies each application in the chart by color and text. You can hover over a section of the chart to view detailed information for that section. Or you can click on an application in the legend to isolate the item in the display and show the same detailed information. You can also click select/deselect an application(s) in the legend to add/remove the application(s) from the display.



Hover the mouse over a section of the pie chart (or click on an application in the legend) to view the number of flows for that application over the time interval displayed (e.g., last 24 hours). In the example below, hovering over the Yahoo section of the pie chart shows the total number of Yahoo flows as 362, or 2.11% of the total number of application flows.



You can also view information on switches passing the application traffic. Click on a section of the pie chart. The pie chart will be broken down by switch for that application. In the example above, clicking on the Yahoo section of the pie chart shows that a single switch (10.255.225.244) is passing 100% (362 flows) of the Yahoo traffic.

List Format

The list format displays a list of applications with flow count information for each one. By default, the list is displayed by application name in alphabetical order; however you can select "Flows" in the **Sort by** drop-down menu to display the applications by flow count. You can also search for and display a specific application by entering the application name in the **Search** field. Click on an application in the table to display switch information for that flow (switch IP address(es) passing the flow, number of flows per switch).

Top N Applications - Advanced App Discovery - List

The screenshot shows the 'Applications' interface with the following data:

Application	Flows
dns	4154
ftp_data	10
google	139
google_ads	68
google_maps	15
Others	14

The interface also features a search bar labeled 'Search for a specific application.' and sorting options: 'Sort by: Name' with a dropdown arrow and 'Sort by name or sample count.' with a search input field.

App Count

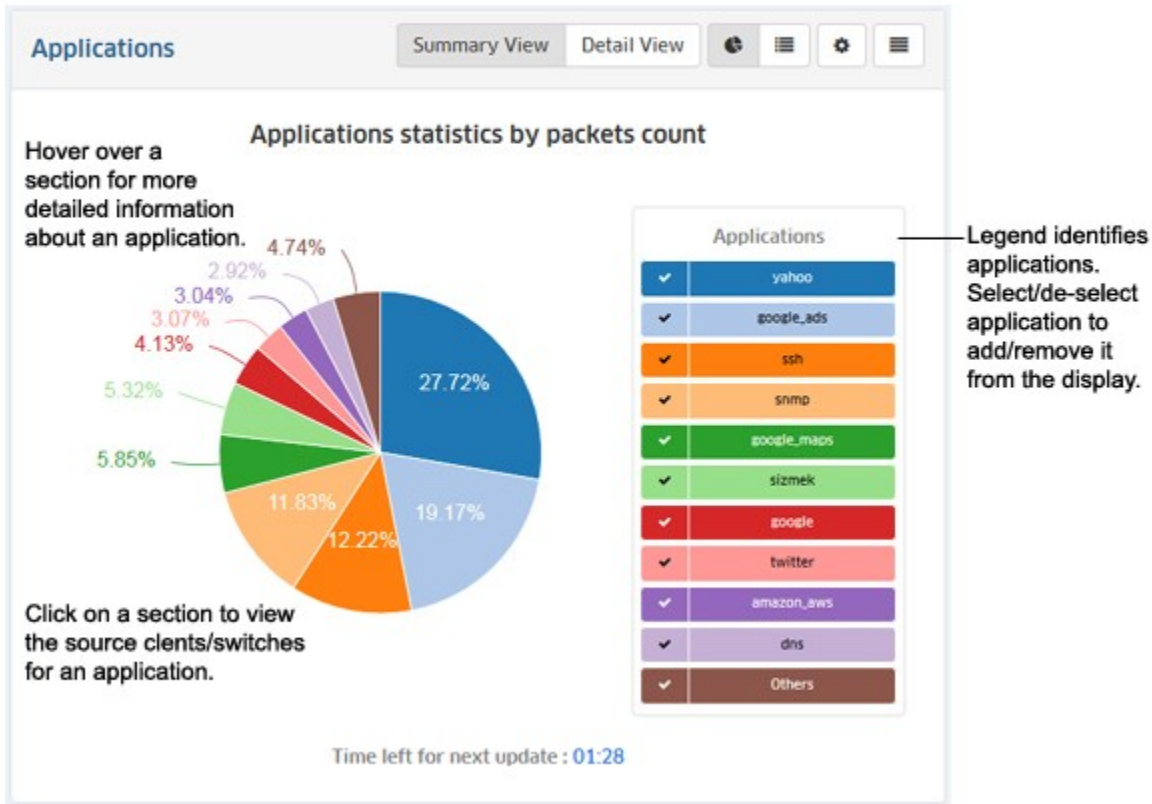
The App Count view displays packet/byte count information for applications/application groups discovered on the network, and the percentage of network resources being used by each application for the selected switches and configured time period. The information can be displayed in a Pie Chart Format (Default) or List Format. Note that App Count is only supported on OS6860/6860E Switches.

Note: In the App Count view, you can view information for all configured applications/groups (Applications), or view applications/groups filtered by configured UNP Profiles (UNP Access Role Profiles). By default, information is displayed for all applications; however, you can display application information by UNP Access Role Profile by clicking on the **UNP Access Role Profiles** button at the top of the screen. The sections below detail the display options for the Application View; however the display options for the UNP Access Role Profile View are similar.

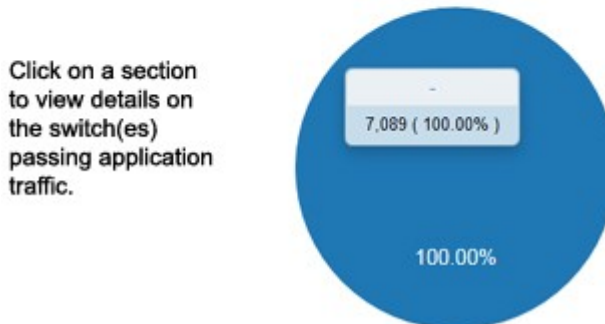
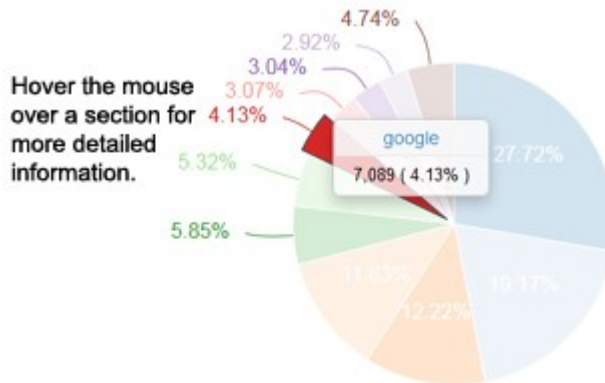
Pie Chart Format

By default, App Count Reports are displayed in Summary View as a pie chart (packet count), with each application displayed as a percentage of the total traffic for the selected switches. By default, an Hourly Summary Report is displayed, showing a summary of the data over the last 24 hours. However, you can configure the report to display different time intervals (e.g., last 24 hours, last 7 days). The legend on the right of the screen identifies each application in the chart by color and text. You can hover over a section of the chart to view detailed information for that section. Or you can click on an application in the legend to isolate the item in the display and show the same detailed information. You can also click select/deselect an application(s) in the legend to add/remove the application(s) from the display. You can also display information in a detailed line graph by clicking on the **Detail View** button at the top of the screen

Top N Applications - Advanced
App Count - Pie Chart



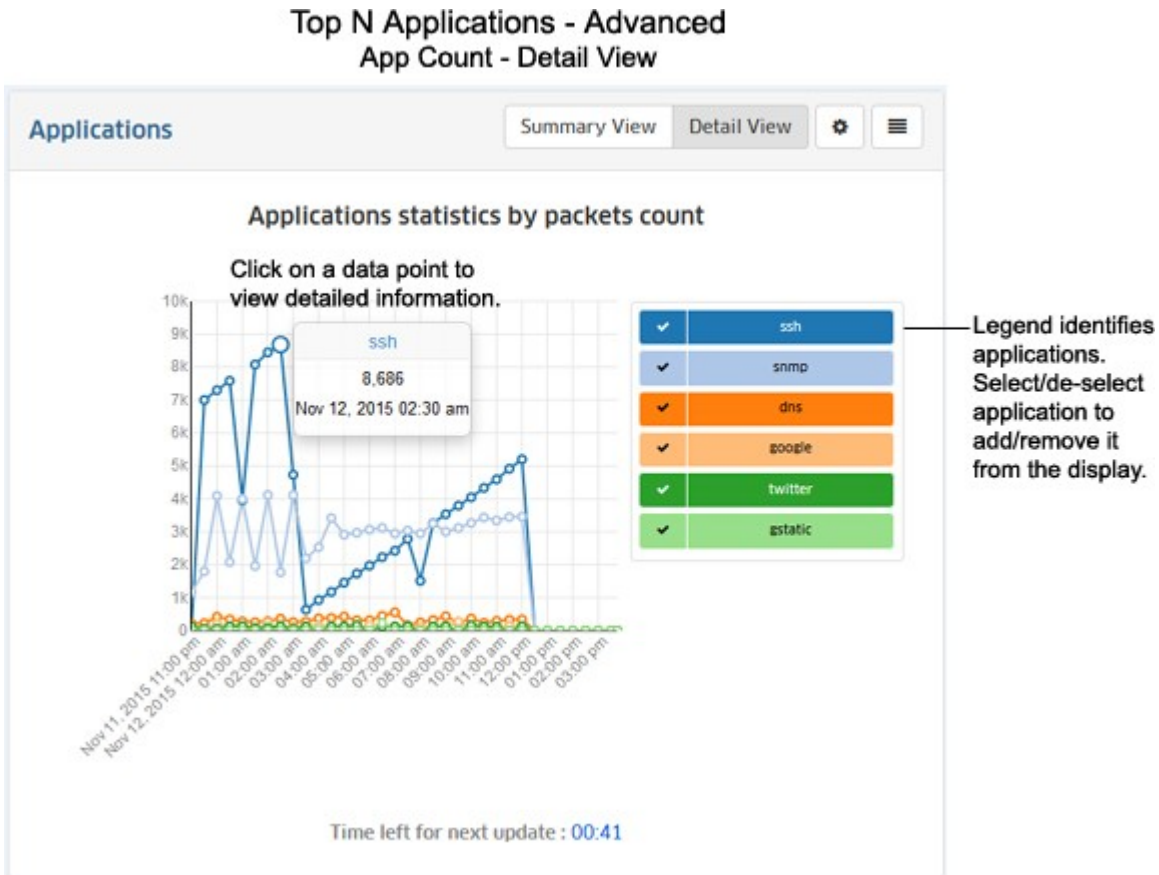
Hover the mouse over a section of the pie chart (or click on an application in the legend) to view the number of packets/bytes for that application over the time interval displayed (e.g., last 24 hours). In the example below, hovering over the Google section of the pie chart shows the total number of Google packets as 7,089, or 4.13% of the total number of application packets.



You can also view information on switches passing the application traffic. Click on a section of the pie chart. The pie chart will be broken down by switch for that application. In the example above, clicking on the Google section of the pie chart shows that a single switch (10.255.225.244) is passing 100% (7,089 packets) of the Google traffic.

Detail View

The Detail View displays detailed data in a line chart for the configured time period. Click on a data point for detailed information. As in the Pie Chart view, the legend on the right of the screen identifies each application in the chart by color and text. And you can also click select/deselect an application(s) in the legend to add/remove the application(s) from the display.



List Format


The list view displays a list of applications with flow count information for each one. By default, the list is displayed by application name in alphabetical order; however you can select "Packet Count/Byte Count " in the **Sort by** drop-down menu to display the applications by flow count. You can also search for and display a specific application by entering the application name in the **Search** field. Click on an application in the table to display switch information for that flow (switch IP address(es) passing the flow, number of flows per switch).

Top N Applications - Advanced App Count - List

The screenshot shows the 'Applications' section of the user interface. At the top, there are tabs for 'Summary View' and 'Detail View', along with several icons for refresh, list, settings, and help. Below the tabs, there is a 'Sort by:' dropdown menu set to 'Name', a 'Sort by name or sample count.' label, and a search input field. A callout box points to the search field with the text 'Search for a specific application.' The main content area displays a list of applications with their respective packet counts:

Application	Packet Count
amazon_aws	5225
dns	5012
google	7089
google_ads	32936
google_maps	10040
Others	8142

Configuring the Information Displayed

You can configure the amount and type of information displayed (e.g., the number of applications displayed, byte or packet information) as well as the time interval that you want to view. To configure the report display, click on the Configuration icon  in the upper-right corner of the report to bring up the Configuration Screen, then complete the fields as described in the following sections. The available options vary depending on the view (e.g., App Discovery, App Count, Packet Count, Byte Count).

- **Choose Chart** - Select the information you want to display:
 - **App Groups** - Displays flow information for all application groups included in the Signature Profile(s) of the selected switch(es) (Default).
 - **Applications** - Displays flow information for all applications included in the Signature Profile(s) of the selected switch(es).
- **Data Interval** - The amount of time, in minutes, between each data point in the Detail View (Range = 15 - 120, Default = 15).
- **Counter Type** - Select whether you want to display data packet or byte count.
- **Data Unit** - Select the date unit for the byte count (Default = MB).
- **Top (apps)** - The number of application groups/applications you want to display (Range = 1 - 50, Default = 5).
- **Time Period Type** - The time interval for the information:
 - **Hourly (Last 24 Hours)** - Displays all information for the last 24 Hours. Use the Time Period field to configure the number of hours of data to display (1 - 24).
 - **Hourly (Any 24 Hours)** - Allows you to display information for a specific 24-hour period over the last week. Use the sliders on the Time Period field to configure the 24-hour period you want to display. You can also set the time period to fewer than 24-hours.

- **Daily (Last 7 Days)** - Allows you to display information for a specific 7-day period over the last 30 days. Use the sliders on the Time Period field to configure the 7-day period you want to display. You can also set the time period to fewer than 7 days.
- **Updating Interval** - How often you want to refresh the data display, in minutes (Range = 1 - 20, Default = 5).

When you are done, click the Apply button. The report display will immediately change to the new view. This will remain the view until it is changed again.

Top N Clients

The Analytics Top N Clients Report Screen displays information for the top network clients including the number of traffic flows for each client. OmniVista uses the sFlow packet to determine the IP address of the client. By default, the Summary View is displayed (pie chart) with each client displayed as a percentage of the total for the configured time interval (e.g., last 24 hours). Information from all network switches in the profile is displayed. However, you can click on the **Select Devices** button to display only information from specific switches. The information can also be displayed in different formats, and you can also configure the amount of information displayed.

To generate a Top N Clients Report, you must first create and assign an Analytics Profile using the Profile Screen that defines the switches/ports that you want to view and the type of information that you want to view on those switches/ports.

Note: sFlow is enabled on a port when you create an Analytics Profile. However, sFlow can also be enabled on a port using the CLI. If sFlow is enabled on a port using the CLI and the OmniVista Server is configured as the receiver, Top N Clients data will be displayed in OmniVista.

Note: sFlow packets cannot be sent through the EMP Port. If you want to gather Top N App data from a switch you cannot use the EMP IP when discovering the switch.

Report Views

The Top N Clients Report can be displayed in a Summary View or a Detail View. The Summary View provides a summary of application traffic for the configured time interval (e.g., last 24 hours (default), last 7 days). The Detail View displays a subset of the data in a bar chart format. For example, if a report is configured to display data for the last 24 hours, the Summary View will display a summary of the data for the last 24 hours; and the Detail View will then display data for each hour within those 24 hours.

Note: Report Views and configuration options are configured using the Options Bar located at the top of the report. This help page contains view and configuration information specific to Top N Clients Reports. For specific information on all of the options available, see the "Report Options" section of the Analytics Reports Help.

Summary View

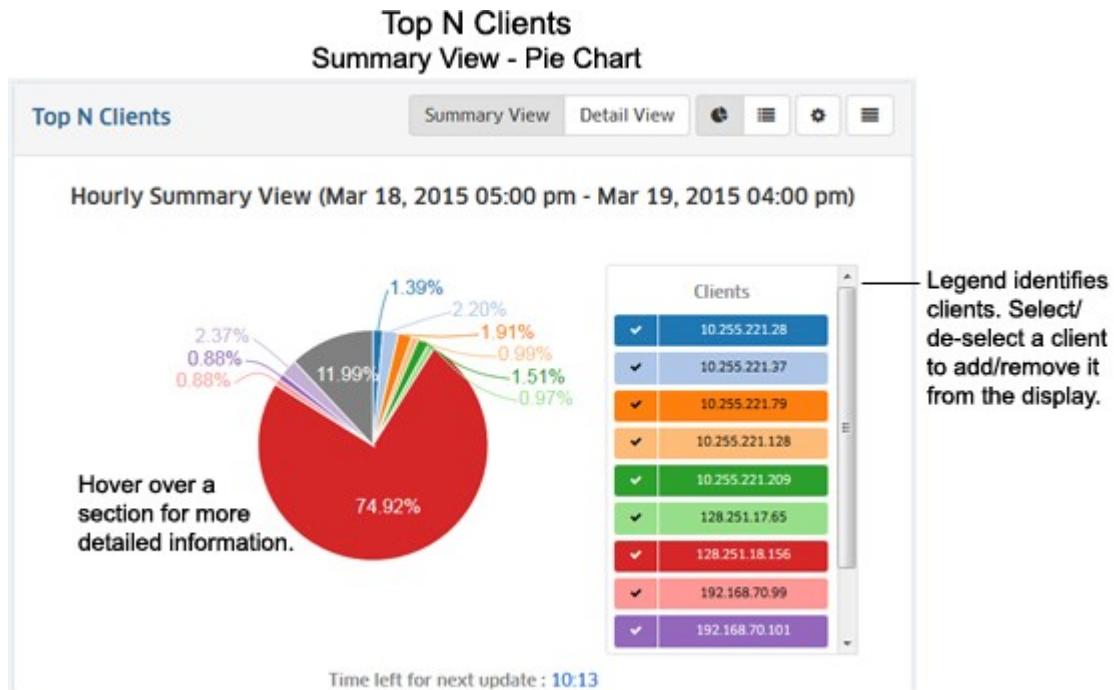
By default, the Summary View is displayed. This view provides a summary of client traffic for the configured time interval (e.g., last 24 hours). By default, the pie chart format is displayed; however a list format is also available.

Pie Chart Format


By default, the Summary View is displayed as a pie chart, with each client displayed as a percentage of the total traffic for all monitored switches for the configured time interval. By default, an Hourly Summary Report is displayed, showing a summary of the data over the last 24 hours. (The Detail View will then display

detailed information for each hour.) However, you can configure the report to display different time intervals (e.g., last 24 hours, last 7 days).

The legend on the right of the screen identifies each client in the chart by color and text. You can hover over a section of the chart to view detailed information for that section. Or you can click on a client in the legend to isolate the item in the display and show the same detailed information. You can also click select/deselect a client(s) in the legend to add/remove a client(s) from the display. The example below shows an Hourly Summary View.



Hover the mouse over a section of the chart (or click on a client in the legend) to view the number of flows for that client over the time interval displayed (e.g., last 24 hours). In the example below, hovering over the section of the pie for client 128.251.18.156, displays the number of flows from that client as 3864, or 74.58% of the total number of traffic flows from that client. You can also view information on applications a client is accessing by clicking on a client section of the pie. The pie chart will change to display all of the applications that the client is accessing. You can also hover over a section of the pie to more detailed information about the application. The legend to the right will identify the applications by color and text. Click on the Back (<) arrow above the legend to return to the previous view.



List Format

The list format displays a list of clients with packet count information for each one. By default, the list is displayed by client IP address in order; however you can select "Samples Count" in the **Sort by** drop-down menu to display the clients by Sample Count. You can also search for and display a specific client by entering the client IP address in the in the **Search** field.

Top N Clients
Summary View - List

The screenshot shows the 'Top N Clients' Summary View interface. At the top, there are tabs for 'Summary View' and 'Detail View'. Below the tabs, the title 'Top N Clients' is displayed. The main content area is titled 'Hourly Summary View (Mar 18, 2015 05:32 pm - Mar 19, 2015 04:32 pm)'. Below this, there is a 'Sort by:' dropdown menu currently set to 'Name', and a 'Search' input field. A callout points to the search field with the text 'Search for specific client.' The list of clients is as follows:

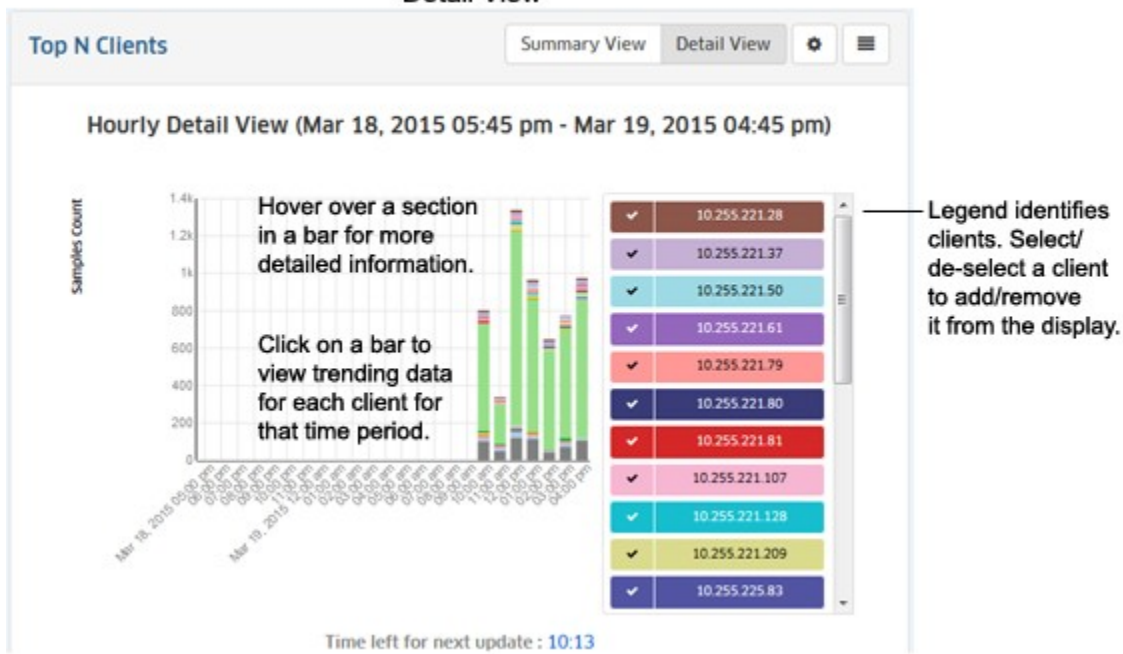
Client IP Address	Samples Count
10.255.221.28	72
10.255.221.37	122
10.255.221.79	104
10.255.221.128	55

Detail View

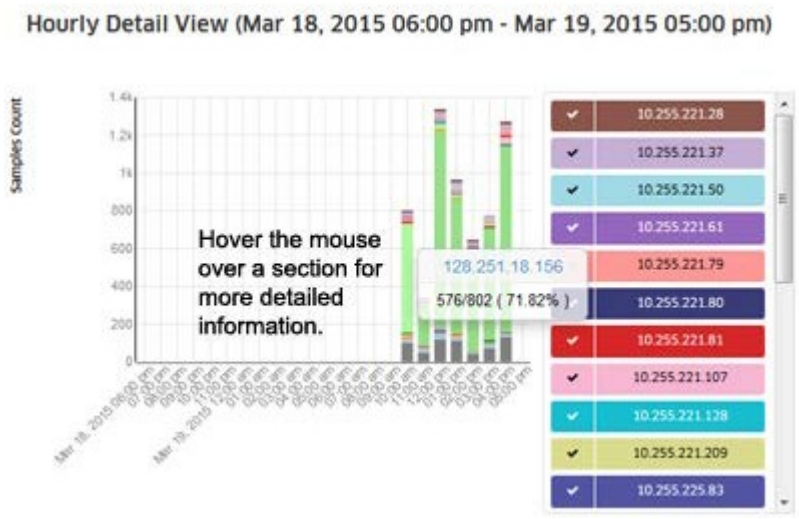
The Detail View displays information in a bar chart view. While the Summary View displays the information for the configured time interval (e.g., last 24 hours), this view provides a detailed view of the specified time interval. For example, if a report is configured to display data for the last 24 hours, the Summary View will display a summary of the data for the last 24 hours; and the Detail View will then display data for each hour within those 24 hours.

Note: You can also click on a bar to view usage trends for that time interval. For example, if you clicked on a day in the chart below, you can view hourly usage trends for each application for that day.

Top N Clients Detail View

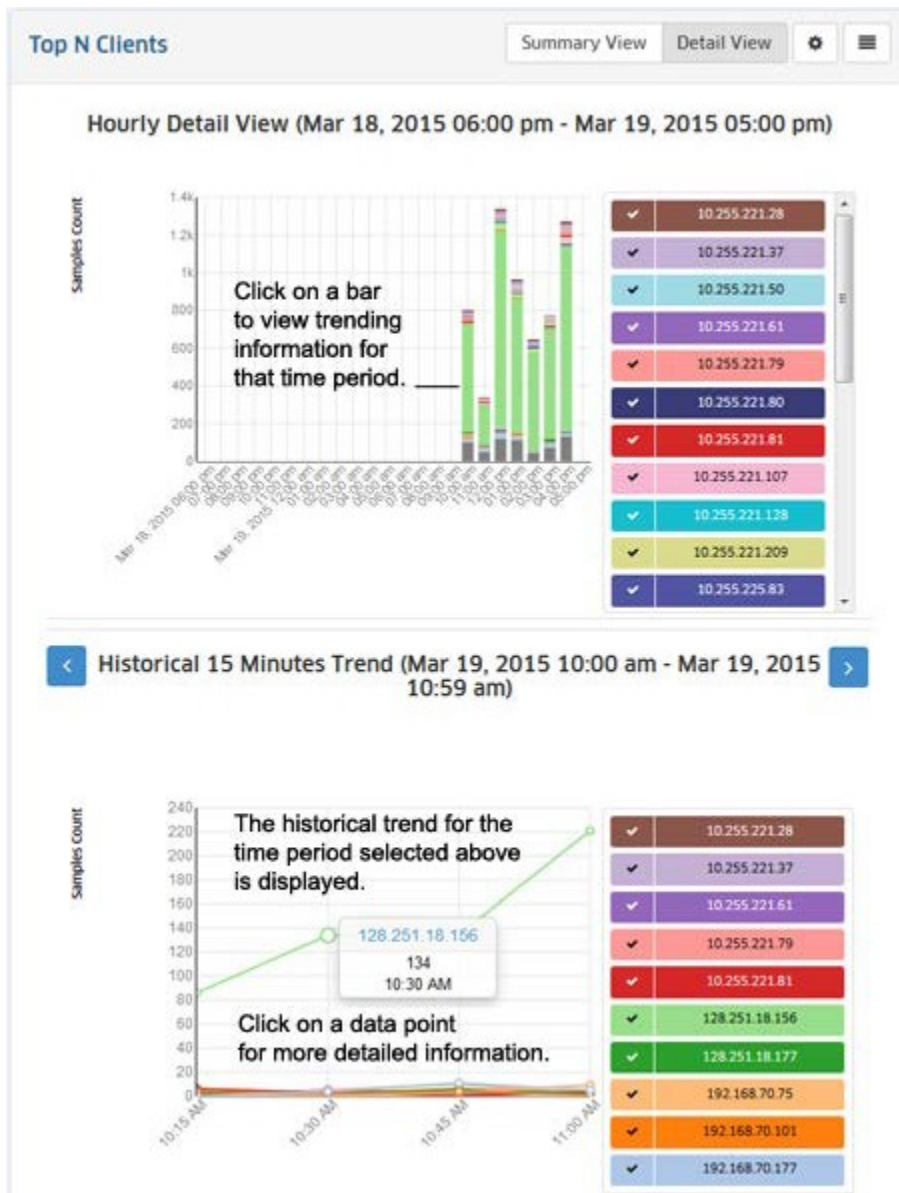


As in the Summary (Pie Chart) view, you can hover the mouse over a section of the chart to view the number of traffic flows for a client over the time interval displayed (e.g., hour). In the example below, hovering over the 128.251.18.156 client of a bar chart shows the total number of traffic flows as 576 out of a total of 802 flows for that hour, or 71.82% of the total number of traffic flows for that hour.




Trending Information

When in Detail View, you can click on a bar in the chart to view usage trends for each client for the selected time interval by "drilling down" on a data set to see a subset of that data. For example, if you selected one of the bars in an Hourly Detail View, the trend for that hour would be displayed in 15 minute increments (as shown below). Click on a data point in the trending view for more detailed information. You can scroll forward or back through the trending date using the arrows at the top of the chart.



Configuring the Information Displayed

You can configure the amount of information displayed (e.g., the number of users you want to view) as well as the time interval that you want to view. To configure the report display, click on the Configuration icon  to bring up the Configuration Screen, then complete the fields as described below to configure how information displayed in the report.

- **Default Devices** - By default, all top switches/ports are displayed. However, you can click on the **Select Devices** button to display only information from specific switches.
- **Number of Top Clients** - The number of clients you want to display (Range = 1 - 20, Default = 5).
- **Interval Type** - The time interval for the information :
 - **Up Until Now** - Displays all information in the selected time interval (e.g., last 24 hours).
 - **Custom** - Set the start and end time for the information you want to display. You can display up to 3 months of data. When data reaches the 3-month maximum, it is overwritten with new data.
- **Time Interval** - The time interval you want to display in the report (e.g. last 24 hours, last 7 days).
- **Auto Refresh Timer** - How often you want to refresh the data display, in minutes (Range = 15 - 60, Default = 15). The configuration option is only available when "Up Until Now" is selected for Interval Type.

When you are done, click the **Save** button. The report display will immediately change to the new view. This will remain the view until it is changed again.

Top N Switches

The Analytics Top N Switches Report Screen displays information for the top switches on the network in terms of the switch's resource usage. Resource usage is calculated based on switch's CPU usage, memory usage, and temperature. If two switches have the same resource usage value, the switch with the highest value in another category will break the tie. By default, the Summary View is displayed (list view). In this view, switches are displayed in a list view from highest to lowest utilization based on the past 24 hours. By default, all monitored switches are displayed. However, you can click on the **Select Devices** button to display only information from specific switches. The information can also be displayed in different formats, and you can also configure the amount of information displayed.

Note: Top N Switches Reports are generated automatically. There is no need to create an Analytics Profile.

Report Views

You can view the Top N Switches Report in a number of ways. The Summary View provides a list of switches displayed from highest to lowest utilization during the configured time interval (e.g., last 24 hours (default), last 7 days). The Detail View displays a subset of the data in a bar chart format. For example, if a report is configured to display data for the last 24 hours, the Summary View will display a summary of the data for the last 24 hours; and the Detail View will then display data for each hour within those 24 hours.

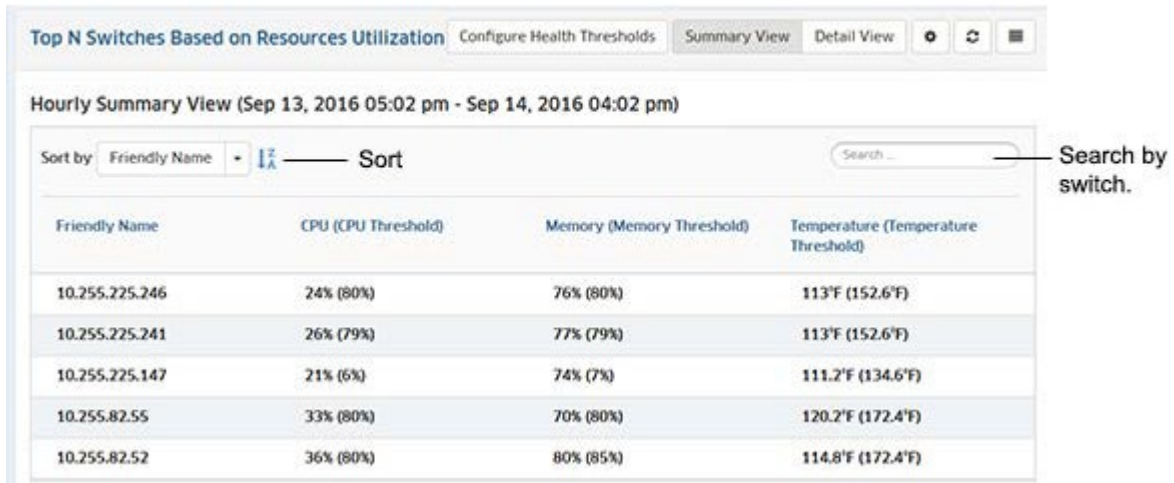
Note: Report Views and configuration options are configured using the Options Bar located at the top of the report. This help page contains view and configuration information specific to Top N Switches Reports. For specific information on all of the options available, see the "Report Options" section of the Analytics Reports Help.

Note: You can also configure health thresholds for any device, by clicking on the **Configure Health Thresholds** button. Health Thresholds are used to set limits for health traps. If a device has been configured to send health traps, a trap will be sent whenever a monitored item's current utilization exceeds the configured health threshold. Configure the CPU, Memory, or Temperature Threshold for the selected device(s) and click on the Apply button. Note that you cannot configure the Temperature Threshold on OS10K, OS6900, or OS6860 devices. The Temperature Threshold is hard coded on these devices. Note that changes made to health thresholds will not appear until the next polling cycle (up to an hour).

Summary View

By default, the Summary View is displayed. In this view, switches are listed from highest to lowest utilization during the configured time interval (e.g., last 24 hours). The usage for each category (CPU, memory, temperature) is displayed as a percentage of the total resources available for each switch.

Top N Switches Summary View

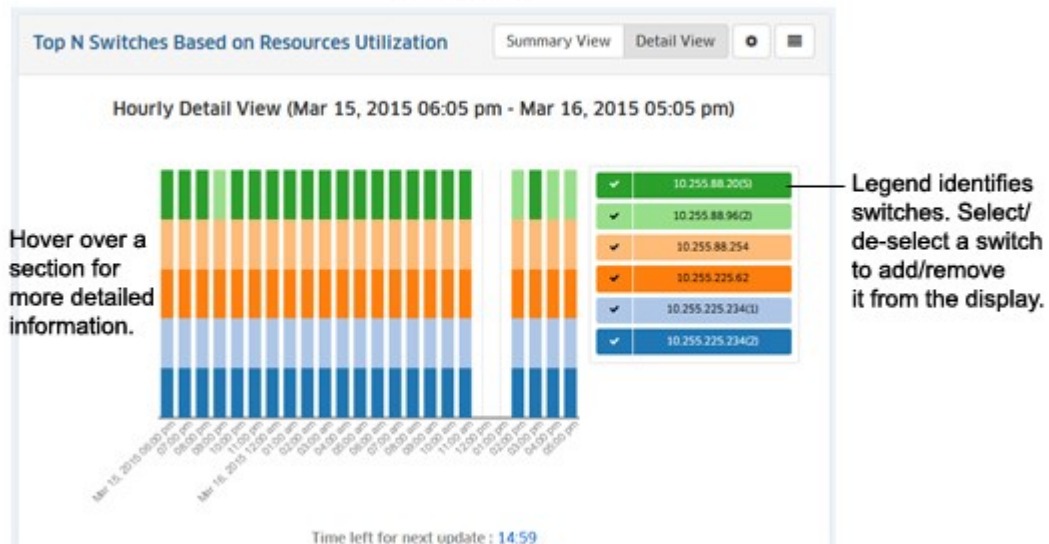


Detail View


The Detail View displays information in a bar chart view. The switches are displayed in order on each bar chart from highest to lowest utilization. While the Summary View displays the information for the configured time interval (e.g., last 24 hours), this view provides a detailed view of the specified time interval. For example, if a report is configured to display data for the last 24 hours, the Detail View will then display data for each hour within those 24 hours. Hover the mouse over a section of the chart (or click on a switch in the legend) to view more detailed information. Click on the Back (<) arrow above the legend to return to the previous view.

The legend on the right of the screen identifies each switch in the chart by color and text. You can hover over a section of the chart to view detailed information for that section. Or you can click on a switch in the legend to isolate the item in the display and show the same detailed information. You can also click select/deselect a switch(es) in the legend to add/remove the switch(es) from the display. The example below shows an Hourly detail View.

Top N Switches Detail View



Configuring the Information Displayed

You can configure the amount of information displayed (e.g., the number of users you want to view) as well as the time interval that you want to view. To configure the report display, click on the Configuration icon  to bring up the Configuration Screen, then complete the fields as described below to configure how information displayed in the report.

- **Default Devices** - By default, all top switches/ports are displayed. However, you can click on the **Select Devices** button to display only information from specific switches.
- **Number of Top Switches** - The number of applications you want to display (Range = 1 - 20, Default = 5).
- **Interval Type** - The time interval for the information :
 - **Up Until Now** - Displays all information in the selected time interval (e.g., last 24 hours).
 - **Custom** - Set the start and end time for the information you want to display. You can display up to 3 months of data. When data reaches the 3-month maximum, it is overwritten with new data.
- **Time Interval** - The time interval you want to display in the report (e.g. last 24 hours, last 7 days).
- **Auto Refresh Timer** - How often you want to refresh the data display, in minutes (Range = 15 - 60, Default = 15). The configuration option is only available when "Up Until Now" is selected for Interval Type.

When you are done, click the **Save** button. The report display will immediately change to the new view. This will remain the view until it is changed again.

Top N Ports Utilization

The Analytics Top N Ports Utilization Report Screen displays the top network ports based on utilization. By default, the Summary View is displayed (list view). In this view, switches/ports are displayed in a list view from highest to lowest utilization for the configured time period (e.g., day, week). By default, all top switches/ports are displayed. However, you can click on the **Select Devices** button to display only information from specific switches. The information can also be displayed in different formats, and you can also configure the amount of information displayed.

Note: To generate a Top N Ports Utilization Report, you must first create an Analytics Profile using the Profile Screen that defines the switches/ports that you want to view and the type of information that you want to view on those switches/ports.

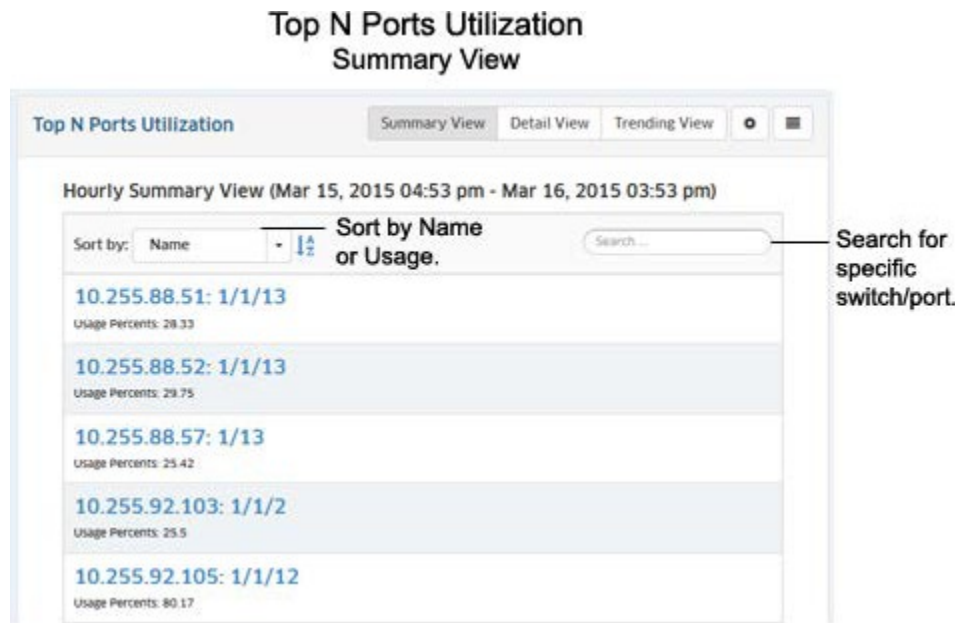
Report Views

The Top N Ports Utilization can be displayed in a Summary View or a Detail View. The Summary View provides a summary of port traffic for the configured time interval (e.g., last 24 hours). The Detail View displays a subset of the data in a bar chart format. For example, if a report is configured to display data for the last 24 hours, the Summary View will display a summary of the data for the last 24 hours; and the Detail View will then display data for each hour within those 24 hours. The Trending View is used to view predicted future port utilization based on past utilization. Port utilization predictions can be used to predict future usage from past trending patterns and provide valuable insight for capacity management.

Note: Report Views and configuration options are configured using the Options Bar located at the top of the report. This help page contains view and configuration information specific to Top N Ports Utilization Reports. For specific information on all of the options available, see the "Report Options" section of the Analytics Reports Help.

Summary View

By default, the Summary View is displayed. In this view, switches/ports are displayed in a list view from highest to lowest utilization for the configured time period (e.g., last 24 hours). Utilization for each port is displayed as a percentage of the total utilization for all monitored ports for the configured time period.

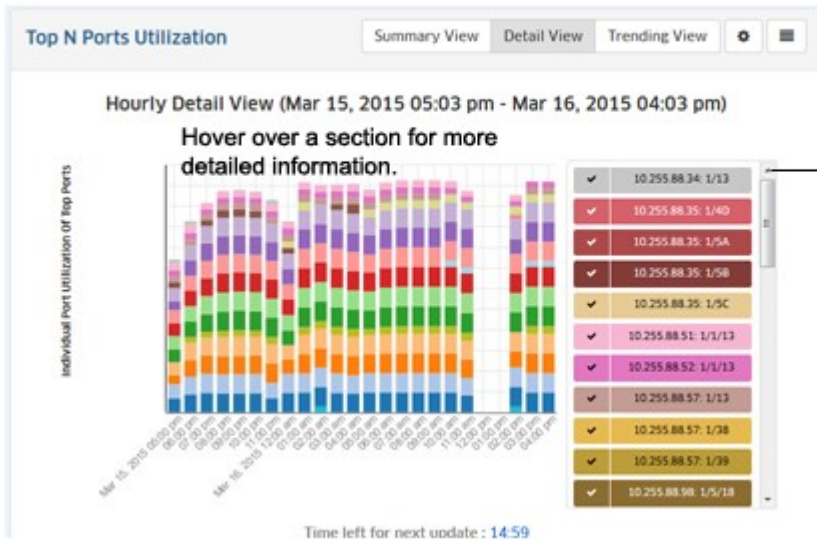


Detail View


The Detail View displays information in a bar chart view. While the Summary View displays the information for the configured time period (e.g., last 24 hours), this view provides a detailed view of the specified time interval. For example, if the Summary View displays information for the last 24 hours, the Detail View will display information for each hour within those 24 hours.

Depending on the number of ports you configured for display (e.g., top 10 ports, top 15 ports), any monitored ports that qualify during the configure time interval (e.g., last 24 hours) are displayed. Ports are simply stacked numerically in each bar by IP address and port number (the order is not based on utilization). The legend on the right of the screen identifies each switch/port in the chart by color and text. You can hover over a section of the chart to view detailed information for that section. Or you can click on a switch/port in the legend to isolate the switch/port in the display and show the same detailed information. You can also click select/deselect a client(s) in the legend to add/remove a client(s) from the display. The example below shows an hourly Detail View for a one day period.

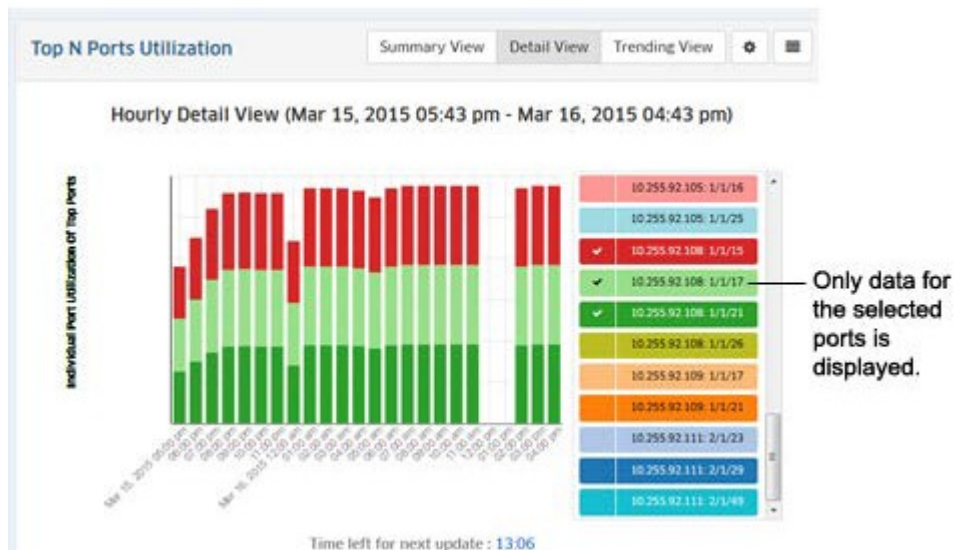
Top N Ports Utilization Detail View



You can hover the mouse over a section of the chart to view the number of flows for that application over the time period displayed (e.g., hour). In the example below, which shows hourly usage, hovering over a section of a bar chart shows switch 10.255.92.106, port 1/1/5 with a utilization of 91% for that hour.



If you want to isolate a port or ports, you can select/deselect the port(s) in the legend.



Trending View

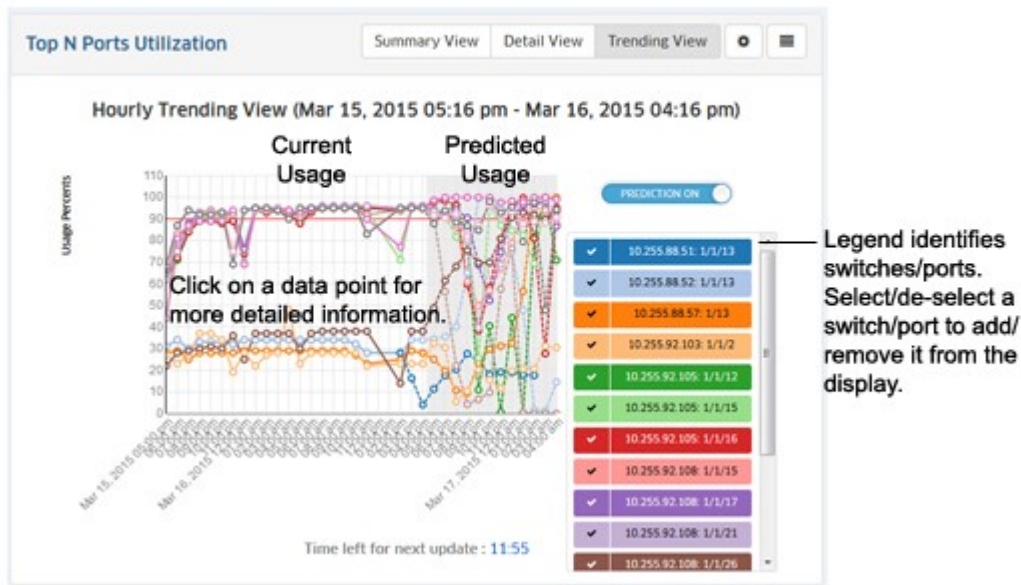
The Trending View is used to view predicted future port utilization based on past utilization. Port utilization predictions can provide valuable insight for capacity management. To make port utilization predictions, OmniVista samples past port utilization for a period of time (Prediction: Training Timeout), and predicts future utilization within a configurable error rate (Prediction: Training Error) using a machine learning algorithm.

To view future trending information, click on the **Trending View** button and enable the **Prediction On** slider. When you initially enable the slider, the slider will display "Prediction in progress" while OmniVista samples and learns port utilization rates. The predicted utilization will then appear in the display to the right of the current utilization. The predicted usage area of the display will be slightly shaded to differentiate it from current usage. The amount of predicted data displayed depends on the interval time configured for the report (e.g., last 24 hours, last 7 days). For predicted data, OmniVista will display approximately one-half of the configured interval time, as shown in the table below.

Configured Time Interval	Amount of Predicted Data
Last 24 Hours	2 Hours
Last 7 Days	3 Days
Last 4 Weeks	2 Weeks


If OmniVista is unable to determine future utilization, a message will appear at the top of the display with a link to the reason(s) (e.g., *10.255.225.234: 1/10. Message: Prediction Analytics for port could not be performed due to insufficient data for Training*).

Top N Ports Utilization
Trending View



The information is displayed in the chart based on the trending configuration settings set in the Trending View Configuration Screen. The screen is also used to set training parameters that OmniVista will use to learn about past/current usage to predict future usage.


Configuring Trending Information

As with other reports, the Trending View Configuration Screen is used to configure how information displayed in the report. It is also used to set training parameters that OmniVista will use to learn about past/current usage to predict future usage. Detailed trending parameters are set in the Preferences Application on the Settings Screen. By default, OmniVista will use these parameters to predict future utilization. Any parameters that you configure on this screen will override the parameters configured in the Preferences Application. Click on the Configuration icon  on the Trending Screen to bring up the Trending Configuration Screen, then complete the fields as described below.

- **Prediction** - Enables/Disables trending prediction.
- **Number of Top Ports** - The number of top ports (in terms of utilization) that you want to display (Range = 1 - 20, Default = 10).
- **Interval Type** - The time interval for the information:
 - **Up Until Now** - Displays all information in the selected time interval (e.g., last 24 Hours).
 - **Custom** - Set the start and end time for the information you want to display. You can display up to 3 months of data. When data reaches the 3-month maximum, it is overwritten with new data.
- **Time Interval** - The time interval you want to display in the report (the past 24 Hours, 7 Days, or 4 Weeks).
- **Auto Refresh Timer** - How often you want to refresh the data display, in minutes (Range = 15 - 60). The configuration option is only available when "Up Until Now" is selected for Interval Type.
- **Threshold** - The threshold level you want to set for the display. A red horizontal line will display on the chart at this threshold level to enable you to quickly see any data that has crossed the level. For example, a threshold of 90, will show a horizontal line at 90% utilization parallel to x-axis of graph.
- **Prediction: Training Timeout** - Specifies how long OmniVista will train, in seconds, by sampling past port utilization. In other words, this specifies how long OmniVista will sample port utilization data before beginning to predict future trends (Range = 15 - 600, Default = 60).
- **Prediction: Training Error** - The target error percentage to which OmniVista will be trained (Default = 0.1 - 1.0, Default = 0.5).

When you are done, click the **Save** button. The report display will immediately change to the new view. This will remain the view until it is changed again.

Configuring the Information Displayed

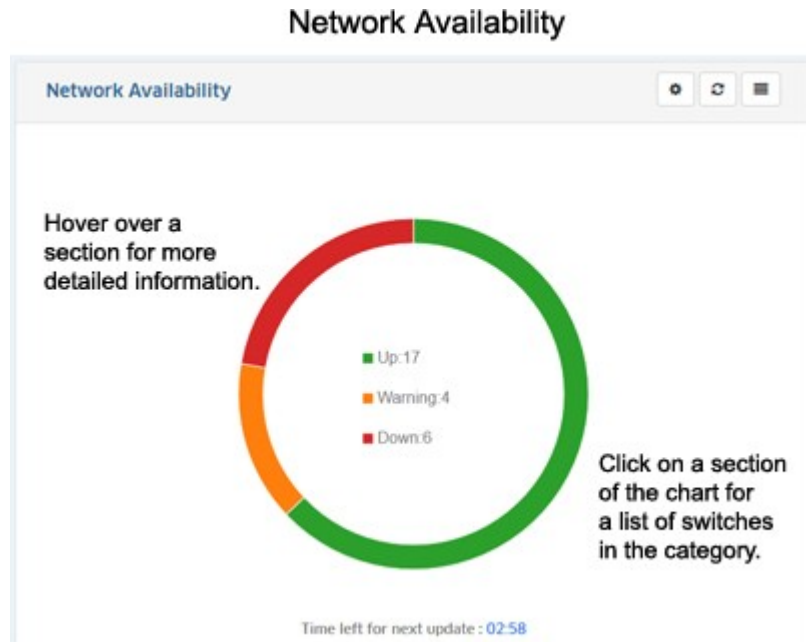
You can configure the amount of information displayed (e.g., the number of ports you want to view) as well as the time period that you want to view. To configure the report display, click on the Configuration icon  to bring up the Configuration Screen, then complete the fields as described below to configure the information displayed in the report.

- **Default Devices** - By default, all top switches/ports are displayed. However, you can click on the **Select Devices** button to display only information from specific switches.
- **Number of Top Ports** - The number of top ports you want to display (Range = 1 - 20, Default = 5).
- **Interval Type** - The time period for the information :
 - **Up Until Now** - Displays all information in the selected time interval (e.g., last 24 hours).
 - **Custom** - Set the start and end time for the information you want to display. You can display up to 3 months of data. When data reaches the 3-month maximum, it is overwritten with new data.
- **Time Interval** - The time interval you want to display in the report (e.g. last 24 hours, last 7 days).
- **Auto Refresh Timer** - How often you want to refresh the data display, in minutes (Range = 15 - 60, Default = 15). The configuration option is only available when "Up Until Now" is selected for Interval Type.

When you are done, click the **Save** button. The report display will immediately change to the new view. This will remain the view until it is changed again.

Network Availability

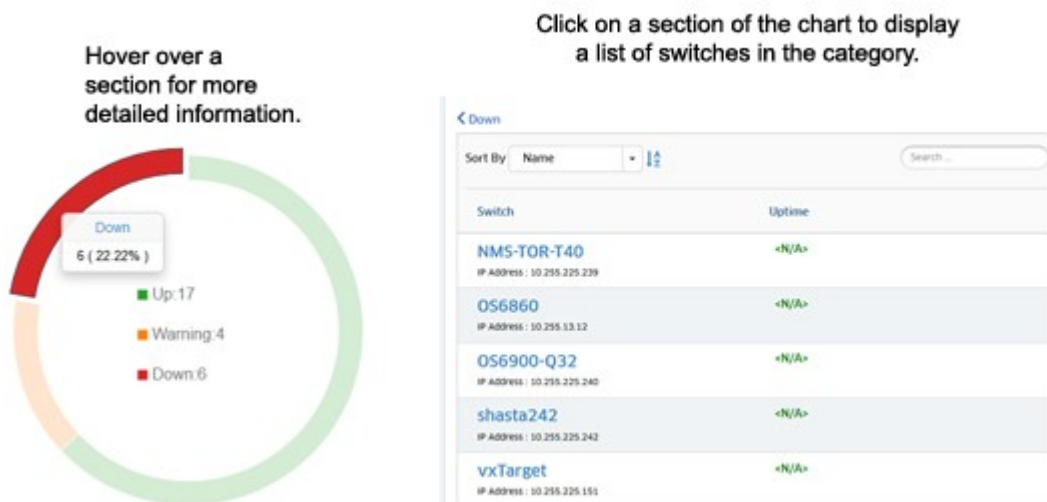
The Analytics Network Availability Screen displays the current operational state of all discovered network devices (Up/Warning/Down). Each category is displayed as a percentage of all monitored switches. The information can be displayed in different formats, and you can configure the information displayed.




Note: Report Views and configuration options are configured using the Options Bar located at the top of the report. This help page contains view and configuration information specific to Network Availability Reports. For specific information on all of the options available, see the "Report Options" section of the Analytics Reports Help.

Report Views

You can view the Network Availability Report in a couple of different ways. Hover the mouse over a category to display a brief summary of the category (the number of switches in the category, along with the percentage of all monitored switches in that category). You can also click on a category to display a list of switches in the category, with specific information about each switch. If you click on a category to display the list view, you can click on the "Back" link (<) to return to the default view.



Configuring the Information Displayed

You can configure the refresh rate for the data displayed by clicking on the Configuration icon  in the Options Bar to bring up the Configuration Screen. Set the Auto Refresh Timer and click on the **Save** button. (Range = 1 - 10 minutes).

Alarms

The Analytics Alarms Screen displays network status/traps for all discovered switches. By default, a graphical pie chart view is displayed. The reported alarms in each severity level are displayed as a percentage of the total alarms reported. You can click on a severity level in the pie chart to view the switch(es) from which the alarms originated, and the number of those alarms received, along with the percentage of the total number of alarms received from that switch. In addition, the information can be displayed in different formats, and you can also configure the amount of information displayed.


Note: Report Views and configuration options are configured using the Options Bar located at the top of the report. This help page contains view and configuration information specific to Alarm Reports. For specific information on all of the options available, see the "Report Options" section of the Analytics Reports Help.

Report Views

You can view the Alarms Report in a number of ways. By default, the pie chart format is displayed. You can also view a list of all alarms.

Pie Chart Format


By default, the pie chart view is displayed. The reported alarms in each severity level are displayed as a percentage of the total alarms reported. You can hover over a section of the chart for more details about the alarm category, or click on a section in the pie chart to view the switch(es) from which the alarms for that category originated.




Hover over a section of the pie chart to display the number of alarms generated for that category. Click on a section of the chart to view information about the switches generating the alarms for that category (the legend will change from Severity Level categories to identify the switches. You can then hover over a section to display detailed information for a specific switch. Click on the Back Arrow (<) above the legend to

return to the default view.

Hover the mouse over a section for more detailed information.



Click on a section (e.g, Normal) to display information about switches in that category.



List Format

The list format displays the exact alarm count for each severity level. By default, the list is displayed by alarm count; however you can select "Severity" in the **Sort by** drop-down menu to display the applications by sample count. You can also search for and display a specific severity level by entering the name in the **Search** field.

Alarms List

Alarms


Sort by: Count Sort by Severity or alarm count. Search: _____

Severity	Count
Normal	599
Warning	401
Critical	0
Major	0
Minor	0

Total: 1 page

Search for specific alarm type.

Configuring the Information Displayed

You can configure the amount of information displayed (e.g., refresh timer). To configure the display, click on the Configuration icon  to bring up the Configuration Screen, then complete the fields as described below to configure the information displayed in the report.

- **Number of Switches** - The number of switches you want to display (Range = 1 - 10).
- **Auto Refresh Timer** - How often you want to refresh the data display, in minutes (Range = 15 - 60).

When you are done, click the **Save** button. The report display will immediately change to the new view. This will remain the view until it is changed again.

Active Calls

The SIP Active Calls Screen is used to display Active Call Record data for selected SIP-enabled switches. To display Active Call Records, select an option from the drop-down menu (Use Switch Picker or Use Topology), then click on the **Select Devices** button to select the switches you want to view. The Active Call Records for the selected switches will be displayed in the table.

By default, the aggregated call records are displayed. The data is an aggregate of all Active Calls on SIP-enabled switches. You can also click on the **View Detailed Call** button at the top of the table to display detailed call records the selected switches.

Viewing Active Call Records

As described above, you can display aggregated or detailed call records in the table. You can also click on a switch(es) in the table to display a graphical representation of the call records.

Aggregated Records

Aggregated Records are call data for any active calls on the network. The data is an aggregate of all active calls on SIP-enabled switches.

- **Device** - The device name.
- **Start Time** - The call start date and time.
- **Calls Count** - The total number of calls processed for SIP Snooping.
- **RTCP Packet Count** - The total number of Real Time Control Protocol (RTCP) packets received by device.
- **RTP Packet Count** - The total number of Real Time Protocol (RTP) packet received by device.
- **Avg Pkt Loss** - The average number of SIP packet received by device.
- **Avg Jitter** - The average jitter, in milliseconds.
- **Avg RTD** - The average Round Trip Delay (RTD).
- **Avg RFactor** - The average RF Facto.
- **Avg MOS** - The average MOS.

Detailed Records

Detailed Records are detailed call data for any active calls on the network. The tab provides detailed data for each Active Call.

- **Device** - The Device name.
- **Call ID** - The call ID.
- **Tag A** - The call tag for call direction A to B.
- **Tag B** - The call tag for call direction B to A.

- **IP Addr A Type** - The IP type for call direction A to B type (e.g., IPv4).
- **IP Addr A** - The IP address for call direction A to B.
- **IP Addr B Type** - The IP type for call direction B to A type (e.g., IPv4).
- **IP Address B** - The IP address for call direction B to A.
- **L4 Port A** - The call L4 port for call direction A to B.
- **L4 Port B** - The call L4 port for call direction B to A.
- **SIP Medial Type** - The SIP Media Type (e.g., Voice, Video)
- **Start Time** - The call start date and time.
- **RTP Count A** - The call Real Time Protocol (RTP) packet count for call direction A to B.
- **RTCP Type A** - The call Real Time Control Protocol (RTCP) packet count for call direction A to B.
- **Rule Name A** - The policy rule name for call direction A to B.
- **RTP Count B** - The call Real Time Protocol (RTP) packet count for call direction B to A.
- **RTCP Count B** - The call Real Time Protocol (RTP) packet count for call direction B to A.
- **Rule Name B** - The policy rule name for call direction B to A.
- **Jitter Violations A** - The call RTCP jitter violations (%) for call direction A to B.
- **Jitter Violations B** - The call RTCP jitter violations (%) for call direction B to A.
- **RTD Violation A** - The call round trip delay violations (%) for call direction A to B.
- **RTD Violation B** - The call round trip delay violations (%) for call direction B to A.
- **Packet Loss Violations A** - Call packet loss violations (%) for call direction A to B.
- **Packet Loss Violations B** - The call packet loss violations (%) for call direction B to A.
- **MOS Violations A** - The call MOS violations (%) for call direction A to B.
- **MOS Violations B** - The call MOS violations (%) for call direction B to A.
- **RF Factor Violations A** - The call RF Factor Violation (%) for call direction A to B.
- **RF Factor Violations B** - The call RF Factor Violation (%) for call direction B to A.
- **Jitter Max A** - The call maximum jitter for call direction A to B.
- **Jitter Min A** - The call minimum jitter for call direction A to B.
- **Jitter Avg A** - The call average jitter for call direction A to B.
- **Jitter Max B** - The call maximum jitter for call direction B to A.
- **Jitter Min B** - The call minimum jitter for call direction B to A.
- **Jitter Avg B** - The call average jitter for call direction B to A.
- **RTD Max A** - The call maximum round trip delay for direction A to B.
- **RTD Min A** - The call minimum round trip delay for direction A to B.
- **RTD Avg A** - The call average round trip delay for direction A to B.
- **RTD Max B** - The call maximum round trip delay for direction B to A.
- **RTD Min B** - The call minimum round trip delay for direction B to A.
- **RTD Avg B** - The call average round trip delay for direction B to A.
- **Pkt Loss Max A** - The call maximum packet loss (%) for call direction A to B.
- **Pkt Loss Min A** - The call minimum packet loss (%) for call direction A to B.
- **Pkt Loss Avg A** - The call average packet loss (%) for call direction A to B.
- **Pkt Loss Max B** - The call maximum packet loss (%) for call direction B to A.
- **Pkt Loss Min B** - The call minimum packet loss (%) for call direction B to A.
- **Pkt Loss Avg B** - The call average packet loss (%) for call direction B to A.
- **RF Factor Max A** - The call maximum RF Factor for call direction A to B.
- **RF Factor Min A** - The call minimum RF Factor for call direction A to B.
- **RF Factor Avg A** - The call average RF Factor for call direction A to B.
- **RF Factor Max B** - The call maximum RF Factor for call direction B to A.
- **RF Factor Min B** - The call minimum RF Factor for call direction B to A.
- **RF Factor Avg B** - The call average RF Factor for call direction B to A.
- **MOS Max A** - The call maximum MOS for call direction A to B.
- **MOS Min A** - The call minimum MOS for call direction A to B.

- **MOS Avg A** - The call average MOS for call direction A to B.
- **MOS Max B** - The call maximum MOS for call direction B to A.
- **MOS Min B** - The call minimum MOS for call direction B to A.
- **MOS Avg B** - The call average MOS for call direction B to A.

Graphical View

You can view a graphical representation of Active Call Records by selecting a switch or switches in the table. By default, the data for "Jitter" is displayed in bar chart format. However, you can select a different variable from the **Variable** drop-down menu; and also change the display to a pie chart by selecting the "Pie" radio button in the **Chart Type** area.

Ended Calls

The SIP Ended Calls Screen is used to display Ended Call Record data for selected SIP-enabled switches. To display Ended Call Records, select an option from the drop-down menu (Use Switch Picker or Use Topology), then click on the **Select Devices** button to select the switches you want to view. You can also configure a Start Time and End Time to only display records from a specific time period.

By default, the aggregated call records are displayed. The data is an aggregate of all Active Calls on SIP-enabled switches. You can also click on the **View Detailed Call** button at the top of the table to display detailed call records the selected switches.

Viewing Ended Call Records

As described above, you can display aggregated or detailed call records in the table. You can also click on a switch(es) in the table to display a graphical representation of the call records.

Aggregated Records

Aggregated Records display call data for any ended calls on the network. The data is an aggregate of all ended calls on SIP-enabled switches.

- **Device** - The device name.
- **Start Time** - The call start date and time.
- **End Time** - The call end date and time.
- **Calls Count** - The total number of calls processed for SIP Snooping.
- **RTCP Packet Count** - The total number of Real Time Control Protocol (RTCP) packets received by device.
- **RTP Packet Count** - The total number of Real Time Protocol (RTP) packet received by device.
- **Avg Pkt Loss** - The average number of SIP packet received by device.
- **Avg Jitter** - The average jitter, in milliseconds.
- **Avg RTD** - The average Round Trip Delay (RTD).
- **Avg RFactor** - The average RF Facto.
- **Avg MOS** - The average MOS.

Detailed Records

Detailed Records are detailed call data for any ended calls on the network. The tab provides detailed data for each ended call. The data is an aggregate of all ended calls.

- **Device** - The Device name.
- **Call ID** - The call ID.

- **Tag A** - The call tag for call direction A to B.
- **Tag B** - The call tag for call direction B to A.
- **IP Address A** - The IP address for call direction A to B.
- **IP Address B** - The IP address for call direction B to A.
- **Port A** - The call L4 port for call direction A to B.
- **Port B** - The call L4 port for call direction B to A.
- **Media Type** - The SIP Media Type (e.g., Voice, Video)
- **Start Date** - The call start date and time.
- **End Date** - The call end date and time.
- **RTP Count A** - The call Real Time Protocol (RTP) packet count for call direction A to B.
- **RTCP Count A** - The call Real Time Control Protocol (RTCP) packet count for call direction A to B.
- **Rule Name A** - The policy rule name for call direction A to B.
- **RTP Type B** - The RTP type for call direction B to A.
- **RTCP Count B** - The call Real Time Protocol (RTP) packet count for call direction B to A.
- **Rule Count B** - The call Real Time Protocol (RTP) packet count for call direction B to A.
- **Rule Name B** - The policy rule name for call direction B to A.
- **End Reason** - The end call reason.
- **Jitter Violation A** - The call RTCP jitter violations (%) for call direction A to B.
- **Jitter Violation B** - The call RTCP jitter violations (%) for call direction B to A.
- **RTD Violation A** - The call round trip delay violations (%) for call direction A to B.
- **RTD Violation B** - The call round trip delay violations (%) for call direction B to A.
- **Packet Loss Violation A** - The call packet loss violations (%) for call direction A to B.
- **Packet Loss Violation B** - The call packet loss violations (%) for call direction B to A.
- **MOS Violation A** - The call MOS violations (%) for call direction A to B.
- **MOS Violation B** - The call MOS violations (%) for call direction B to A.
- **RF Factor Violation A** - The call RF Factor Violation (%) for call direction A to B.
- **RF Factor Violation B** - The call RF Factor Violation (%) for call direction B to A.
- **Jitter Max A** - The call maximum jitter for call direction A to B.
- **Jitter Min A** - The call minimum jitter for call direction A to B.
- **Jitter Avg A** - The call average jitter for call direction A to B.
- **Jitter Max B** - The call maximum jitter for call direction B to A.
- **Jitter Min B** - The call minimum jitter for call direction B to A.
- **Jitter Avg B** - The call average jitter for call direction B to A.
- **RTD Max A** - The call maximum round trip delay for direction A to B.
- **RTD Min A** - The call minimum round trip delay for direction A to B.
- **RTD Avg A** - The call average round trip delay for direction A to B.
- **RTD Max B** - The call maximum round trip delay for direction B to A.
- **RTD Min B** - The call minimum round trip delay for direction B to A.
- **RTD Avg B** - The call average round trip delay for direction B to A.
- **Packet Loss Max A** - The call maximum packet loss (%) for call direction A to B.
- **Packet Loss Min A** - The call minimum packet loss (%) for call direction A to B.
- **Packet Loss Avg A** - The call average packet loss (%) for call direction A to B.
- **Packet Loss Max B** - The call maximum packet loss (%) for call direction B to A.
- **Packet Loss Min B** - The call minimum packet loss (%) for call direction B to A.
- **Packet Loss Avg B** - The call average packet loss (%) for call direction B to A.
- **RF Factor Max A** - The call maximum RF Factor for call direction A to B.
- **RF Factor Min A** - The call minimum RF Factor for call direction A to B.
- **RF Factor Avg A** - The call average RF Factor for call direction A to B.
- **RF Factor Max B** - The call maximum RF Factor for call direction B to A.
- **RF Factor Min B** - The call minimum RF Factor for call direction B to A.
- **RF Factor Avg B** - The call average RF Factor for call direction B to A.

- **MOS Max A** - The call maximum MOS for call direction A to B.
- **MOS Min A** - The call minimum MOS for call direction A to B.
- **MOS Avg A** - The call average MOS for call direction A to B.
- **MOS Max B** - The call maximum MOS for call direction B to A.
- **MOS Min B** - The call minimum MOS for call direction B to A.
- **MOS Avg B** - The call average MOS for call direction B to A.

Graphical View

You can view a graphical representation of Active Call Records by selecting a switch or switches in the table. By default, the data for "Jitter" is displayed in bar chart format. However, you can select a different variable from the **Variable** drop-down menu; and also change the display to a pie chart by selecting the "Pie" radio button in the **Chart Type** area.

Performance Monitoring

The Analytics Performance Monitoring Feature enables you to collect, monitor, and view performance statistics for devices throughout the network. You can create customized line graphs to show device, module, and port health, and store that data on the OmniVista Server. You can also display the data in Table format for a more detailed view. The Performance Monitoring feature performs its own independent polling to collect data, and this polling can be toggled on and off when desired. The polling rate can also be configured.

The Performance Monitoring feature is configured by creating Statistics Profiles that specify the devices and variables (e.g., CPU usage, port utilization) that you want to monitor. The Performance Monitoring Screen (below) displays all configured Statistics Profiles and is used to create, edit, and delete profiles. It is also used to view profile data, schedule profiles, and start and stop profiles.

Performance Monitoring

<input type="checkbox"/>	Name	Status	Running Mode	Start Time	End Time	Interval
<input type="checkbox"/>	Switch 1	RUNNING	MANUAL	May 15, 2017 3:16:09 pm		60
<input type="checkbox"/>	Switch 2	RUNNING	MANUAL	May 15, 2017 3:17:59 pm		60
<input type="checkbox"/>	Switch 3	RUNNING	MANUAL	May 16, 2017 10:48:13 am		60

Showing All 3 rows

The following statistics can be generated for network devices:

- Switch Health (Overall Device Resource Utilization)
 - Rx Utilization
 - RxTx Utilization
 - Memory Utilization
 - CPU Utilization
 - Switch Temperature
- Module Health (Specific NI Utilization)
 - Rx Utilization
 - RxTx Utilization
 - Memory Utilization
 - CPU Utilization
- Port Health (Specific Port Utilization)
 - Rx Utilization
 - RxTx Utilization
- Ethernet Ports (Specific Port Traffic Information)
 - Port Rx CRC Error Frames
 - Port Rx Lost Frames
 - Port Rx Error Frames
 - Port Tx Lost Frames
 - Port Tx Collided Frames
 - Port Tx Error Frames
 - Port Rx Alignment.

Note: Not all variables are supported on all device types. When you are creating a profile, only supported variables for the device type are displayed for selection.

Using the Performance Monitoring Feature

The Performance Monitoring feature is configured by creating Statistics Profiles that specify the devices and variables (e.g., CPU usage, port utilization) that you want to monitor, as well as the polling intervals, data retention, and graphing attributes you want to use to display the data (e.g., line weights/colors). Profiles are created on the Performance Monitoring Screen and displayed in the Statistics Profile List (shown above). The list provides basic profile information and is used to create, edit, and delete, profiles, as well as view profile data. It also used to stop and start profile data collection.

Creating a Statistics Profile

To start collecting and displaying data, you must create a Statistics Profile containing the devices and variables that you want to monitor. Click on the Add icon **+** at the top of the Performance Monitoring Screen to open the Create Profile Wizard, and complete the screens as described below to create the profile.

Statistics Properties Screen

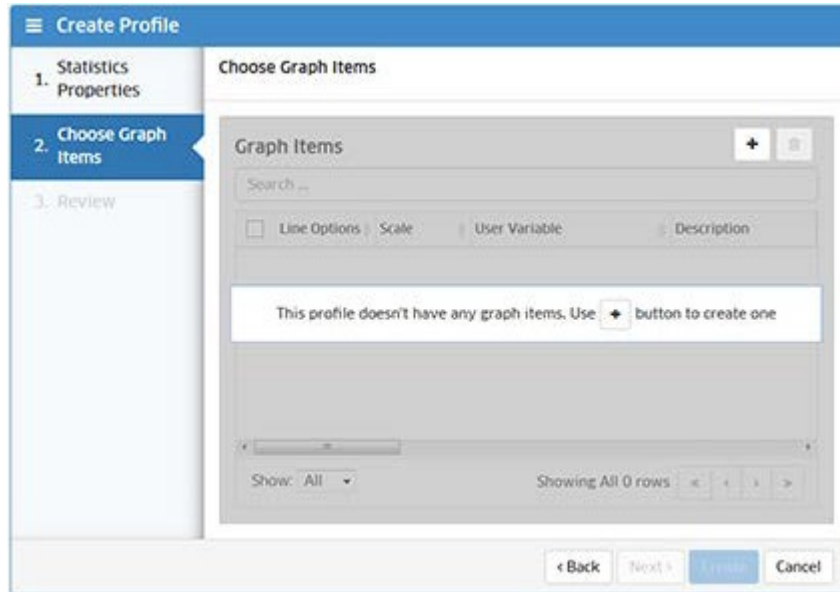
The Statistics Properties Screen in the wizard is used to create a name/description for the profile and set monitoring parameters. Complete the fields as described below and click **Next** to go to the Choose Graph Items Screen.

- **Profile Name** - The profile name. It is advisable to use meaningful names and/or descriptions for profiles, so that when glancing at the Statistics Profile list you will be able to distinguish one from another. For example, you may want to add a device name or IP address, and the type of variable you are monitoring, such as CPU or port utilization.
- **Description** - An optional profile description.
- **Auto Poll** - Enables (On)/Disables (Off) automatic polling for the profile. OmniVista will continuously poll the devices in the profile based on the polling interval and Data Retention Period. Disabling automatic polling performs the same function as stopping a profile on the Statistics Profile List Screen. There may be times when you want to temporarily disable automatic polling. For example, if you wish to spend more time viewing and analyzing a certain group of statistics, automatic polling will interfere by updating these statistics when the next polling cycle is performed. To prevent this, you can edit the profile to disable automatic polling (Default = On).
- **Interval** - When automatic polling is enabled, specify the desired polling rate, in seconds (Range = 20 - 60, Default = 60).
- **Data Retention Period** - The amount of time Statistics Data is saved on the OmniVista Server. You can specify the retention period in Days (1 - 180) or Hours (1 - 24). When the specified data retention period is reached, new incoming data will overwrite the oldest data. For example, if you specify a data collection size of 30 days, new incoming data will begin to overwrite the oldest data after 30 days of data have been saved to the server. (Range = 1 - 180 Days, Default = 30 Days)

Note: By default, the Advanced Options are not displayed when the Statistics Properties Window is displayed. Click on the "Show Advanced Options" link to display the Advanced Options Fields (Auto Poll, Interval, and Data Retention Period).

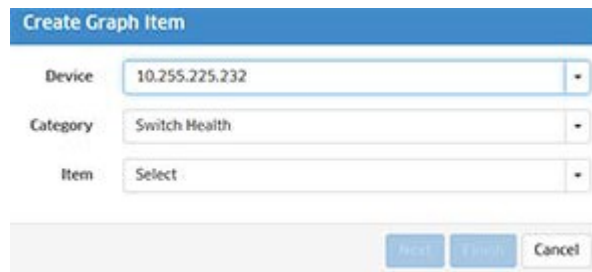
Choose Graph Items Screen

The Choose Graph Items Screen in the wizard is used to select the devices/variables (Counters) you want to monitor. It is also used to customize the way the data is displayed.



When you are creating a new profile, there are no Graph Items displayed, as shown above. Click on the Add icon **+** to bring up the Create Graph Item Window (shown below) to select the device(s) and variables that you want to monitor. Select the device(s) and variables as described below. After selecting a device/variable, click on the **Next** button to add additional devices/variables to the profile. Repeat to add additional devices/variables. When you have included all of the devices/variables you want in the profile. Click on the **Finish** button.

Note: Note that you can have a maximum of 50 Counters in a single profile.

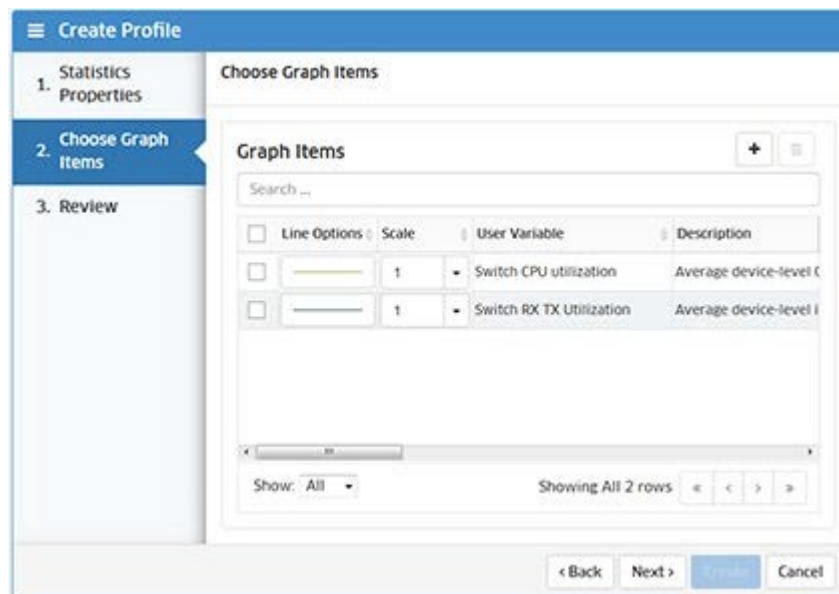


- **Device** - Select a device that you want to include in the profile.
- **Category** - Select a variable category. Note that not all variables are supported on all device types. When a device is selected, only supported variables are displayed in the Category and Item drop-down menus.
 - **Switch Health** - Select one of the following variables for a switch:
 - **Rx Utilization** - Percentage of switch Rx capacity being used
 - **Rx/Tx Utilization** - Percentage of switch Rx/Tx capacity being used
 - **Switch Memory Utilization** - Percentage of switch memory capacity being used
 - **CPU Utilization** - Percentage of switch CPU capacity being used
 - **Switch Temperature** - Switch temperature, degrees Fahrenheit.
 - **Module Health** - Select one of the following variables for a module:
 - **Rx Utilization** - Percentage of module Rx capacity being used
 - **RxTx Utilization** - Percentage of module RxTx capacity being used
 - **Memory Utilization** - Percentage of module memory capacity being used
 - **CPU Utilization** - Percentage of module CPU capacity being used.
 - **Port Health** - Select one of the following variables for a port:
 - **Rx Utilization** - Percentage of port Rx capacity being used
 - **RxTx Utilization** - Percentage of port RxTx capacity being used.

- **Ethernet Ports** - Select one of the following variables for a port:
 - **Port Rx CRC Error Frames** - Number of CRC error frames received on the port
 - **Port Rx Lost Frames** - Number of lost frames received on the port
 - **Port Rx Error Frames** - Number of error frames received on the port
 - **Port Tx Lost Frames** - Number of lost frames transmitted on the port
 - **Port Tx Collided Frames** - Number of collision frames received or transmitted
 - **Port Tx Error Frames** - Number of error frames transmitted on the port
 - **Port Rx Alignment** - Number of alignments error frames received on the port.

You can only select one group of variables at a time (Device, Category, Item). For example, you could select Device "10.255.225.252", Category "Switch Health", and Item "Switch Memory Utilization". After selecting these variables you could click on the **Next** button to add additional variables for the device, or add new devices/variables to the profile; or just click on the **Finish** button to create the profile. When you click on the **Finish** button, the graph items you selected are displayed.

In the example below, two variables were selected for the profile - "Switch CPU Utilization" and "Switch Rx/Tx Utilization". You can click on the Add icon **+** to add additional devices/variables, or select a variable and click on the Delete icon **✖** to remove it from the profile.



Each variable is automatically assigned a line color and default width. You can change the line color or width of a variable's line by clicking on the line in the Line Options Column and selecting a different color and/or width. You can also change the scale used to plot a variable by clicking on Scale Column and selecting a different scale. All of the items in each variable are displayed in the Graph Items Table, as described below.

- **Line Options** - The line color and thickness used to display the variable in the line chart.
- **Scale** - The scale used on the line chart for the variable. By default, all variables (except Ethernet Tx/Rx Bytes) are set to a scale of 1. By default, Ethernet Tx/Rx Bytes are set to a scale of 0.000001. When two or more variables are being graphed simultaneously, the graph line for one variable may be disproportionately large (or small), making it difficult to view. If this situation occurs, you may wish to select a different scale for that variable. In this context, scale is really a multiplier. The range of scale values is 1,000,000 - 0.000001.
- **User Variable** - The specific user variable being monitored.
- **IP Address** - The IP address of the device being monitored.
- **Chassis** - The number of the chassis being monitored, if applicable.
- **Slot** - The slot being monitored, if applicable.

- **Port** - The port being monitored, if applicable
- **Device Name** - The name of the device being monitored.
- **Device Type** - The specific device type being monitored (e.g., OS6850-48X).
- **Category** - The category of the variable (e.g., Switch Health, Module Health, Port Health).
- **Description** - A brief description of the variable being monitored.

Review Screen

The Review Screen enables you to review the profile configuration. If necessary, click on the **Back** button to go to a previous screen and modify the profile. Click on the **Create** button, then click on the **Finish** button to create the profile. The profile will begin collecting data.

Create Profile

Review

Profile Name: Switch 1
 Description: Switch Health
 Auto Poll: ON
 Interval: 60
 Data Retention Period: 30 day(s)

Graph Items

Line Options	Scale	User Variable	Description
	1	Switch CPU Utilization	Average device-level C
	1	Port TxRx Utilization	Average port-level I/O
	1	Port TxRx Utilization	Average port-level I/O
	1	Switch CPU Utilization	Average device-level C
	1	Switch CPU Utilization	Average device-level C

Show: All Showing All 6 rows

< Back Next > **Create** Cancel

Note: After creating a profile you can always edit the profile to add/remove variables, or modify line options or scale.

Viewing Statistics Profile Data

To view profile data, select the profile in the Profile Statistics List and click on the **View Statistics** button at the top of the list. You can only view data for one profile at a time. By default, data is displayed for the last hour. However, you can change the display time range or display more detail within a time range. You can also change the line color/width or scale of a variable to make it easier to view in the chart; and you can add or remove variables to/from a profile.

Note: You can click on the **Switch to Table** button at the top of the screen to display detailed data in table view. From Table View, you can click on the **View Chart** button to return to chart view.

Switch 1 Statistics [last 1 hour](#) — Click to change the display time range.

Poll Save to PNG Add Counter

Counters

<input type="checkbox"/>	Line Options	Scale	User Variable
<input type="checkbox"/>		1	Switch CPU Utilization
<input type="checkbox"/>		1	Port TdRx Utilization
<input type="checkbox"/>		1	Port TdRx Utilization
<input type="checkbox"/>		1	Switch CPU Utilization
<input type="checkbox"/>		1	Switch CPU Utilization

Show: All Showing All 6 rows

Expand Time Range

Scroll through data

Click to add variables to the profile.

Select a variable(s) and click on **Show** or **Hide** to add or remove them from the display.

As shown in the example above, the data for the profile (e.g., Switch 1 Statistics) is displayed in a line chart on the left side of the screen and the variables included in the profile are displayed on the right side of the screen in the Counters Table. You can use the Counters Table as a legend to view different variables in the line chart. You can also select a variable(s) and click on the **Show** or **Hide** buttons to add or remove variables from the chart. When you hide a variable, the variable will be "grayed out" in the Counters Table. Data will still be collected for the variable(s), the data will just not be displayed in the chart. Select the hidden variable(s) and click on the **Show** button to display the variable in the chart.

Note: The Performance Monitoring Feature automatically polls devices in a profile based on the Automatic Polling Interval configured for the profile. However, you can click on the **Poll** button above the chart to perform an immediate poll and re-display of the data. You can also click on the **Save to PNG** button to save/view the chart as a .png file.

Changing the Display Time Range

By default, data is displayed for the last hour. Click on the time range link (e.g., "last 1 hour") to bring up the Time Range Window (shown below). Change the time range as described below and click on the **Get Data** button.

Time Range Until Now Custom

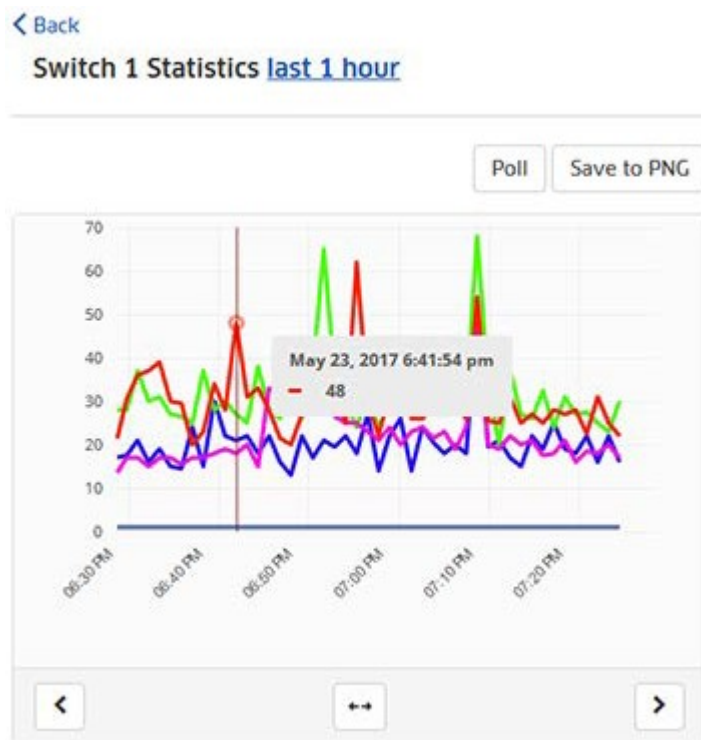
Duration All Data (30 days)
 Last 7 days
 Last 24 hours
 Last 1 hour

- **Until Now** - Displays data from the last time the profile was started until the present time. If you select this option, you can set a duration:
 - **All Data** - Displays all data collected during the Data Retention Period configured for the profile (e.g., 30 days).
 - **Last 7 Days** - Displays data for the last 7 days.
 - **Last 24 Hours** - Displays data for the last 24 hours.
 - **Last Hour** - Displays data for the last hour.
- **Custom** - Set a Start Time and End Time for the data display. Only data collected during this time period is displayed.

Note: If you increase the time range on a profile, (e.g., go from "Last Hour" to "Last 24 Hours"), OmniVista displays the average number for each data point. For example, when viewing data over the last hour, you may see significant data variations in a variable when viewing it over a given hour. But when you change to a 24-hour view, OmniVista will display the data for that hour as the average for all data collected during that hour. You can always select a time period to "zoom in" on a specific time for a more detailed view of the data.

Displaying More Detail in a Time Range

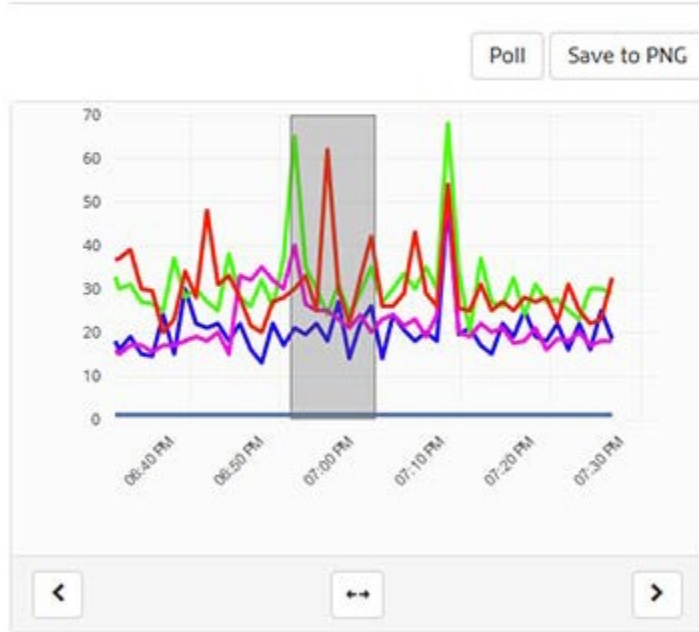
You can place the cursor on a variable line anywhere along the timeline to view exact data for that variable at a specific time, as shown below. The data displayed is based on the polling interval configured for the profile. For example, if the polling interval is 60 seconds, data is displayed at one minute intervals as you move the cursor along the timeline. The line color for the variable is displayed along with the data. Specific data for each poll is also available in Table View.



You can also zoom in on a specific period on a time range to view more detail. Click and drag the mouse along the section of the graph that you want to view. The section is highlighted in gray, as shown below.

[Back](#)

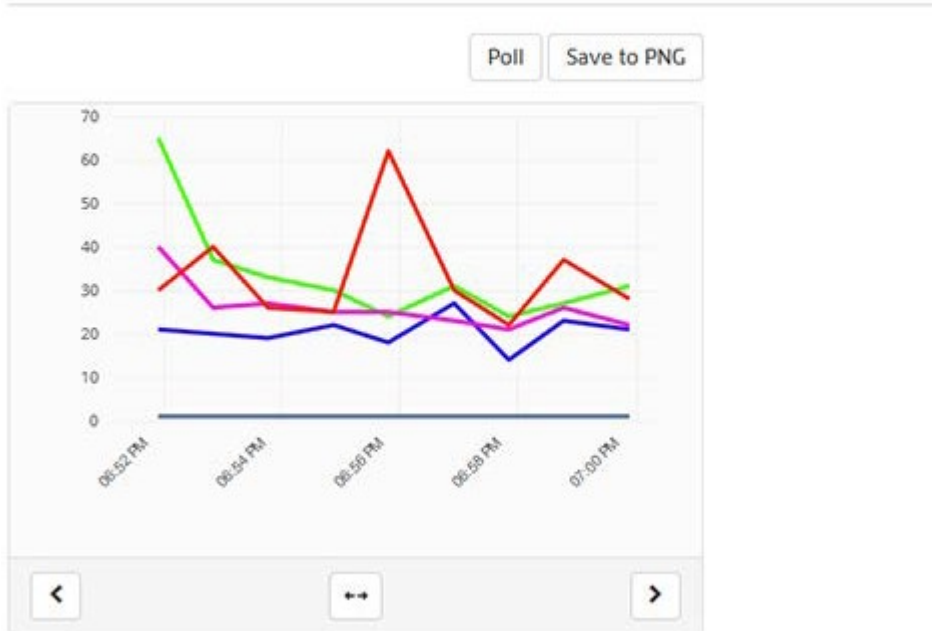
Switch 1 Statistics [last 1 hour](#)



Release the mouse clicker and the time range you selected will be displayed in detail. As you can see, rather than displaying the last hour of data, the graph now displays a detailed view of the selected time range (May 23, 2017 6:51:08 p.m. - May 23, 2017 7:00:18 p.m.). Click on the time range link to return the view to "last hour".

[Back](#)

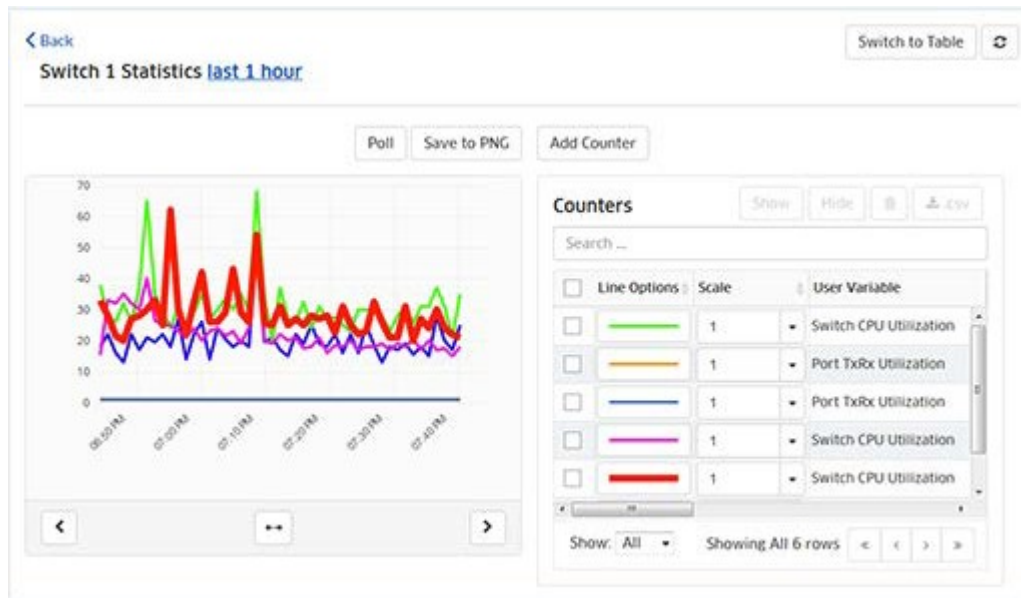
Switch 1 Statistics [May 23, 2017 6:51:08 pm - May 23, 2017 7:00:18 pm](#)



Changing A Variable's Line or Scale

You can change the color or width of a variable's line to help you track the variable more easily. To change a line color or width, click on the line in the Line Options Column of the Counters Table. To change the line color, click on the Line Color drop-down, choose a different line color, then click on the **Choose** button. To


change a line width, click on the Line Width drop-down and select a different width. When you are done, click anywhere in the display to close the window. The variable will be displayed in the chart with the new color and/or width. In the example below, the line width of the Switch CPU Utilization variable at the bottom of the Counters Table was changed to make it stand out from the other variables in the chart.



You can also change the scale used to plot a variable by clicking on Scale Column and selecting a different scale. By default, all variables are set to a Scale of 1. When two or more variables are being graphed simultaneously, the graph line for one variable may be disproportionately large (or small), making it difficult to view. If this situation occurs, you can select a different scale for that variable. In this context, scale is really a multiplier. The range of scale values is 1,000,000 - 0.000001.


Adding/Deleting Variables

To add a variable to the profile, click on the **Add Counter** button at the top of the Counters Table and add a variable (Device, Category, Item) as described above. When you are done adding the variable, click on the **Add** button. Repeat to add additional variables. When you are done, click on the **Done** button. The variable(s) will be added to the profile and displayed at the top of the Counters Table, and data will start being collected for the new variable(s).

You can also select a variable in the Counters Table and click on the Delete icon  to remove the variable from the profile. This will remove the variable from the profile, not just the display.

Note: Note that you can have a maximum of 50 Counters in a single profile.

Table View

To display data in a detailed table view, click on the **Switch to Table** button at the top of a profile chart screen. You can only view statistics for one device at a time. Click on the **Select a Device** button, select a device, then click on **OK** to display the data. Table View provides detailed information at each data poll. The information is displayed as a percentage of the maximum value for each variable. For example, in the table below, the most recent poll shows that the selected device used 22 percent of its maximum CPU capacity, 73 percent of its Switch Memory capacity, and 1 percent of the its TxRx capacity on port 1/47. Click on the **Select a Device** button to view data for a different device in the profile. Click on the **View Chart** button to return to chart view. You can also click on the  **.CSV** button at the top of the table to view/save the table as a .csv file.

Time	Switch CPU utilization	Switch Memory Utilization	Port TxRx Utilization (Port 1/47)	Port TxRx Utilization (Port 1/46)
May 16, 2017 4:32:31 pm	22.00	73.00	1.00	1.00
May 16, 2017 4:31:31 pm	17.00	73.00	1.00	1.00
May 16, 2017 4:30:31 pm	24.00	73.00	1.00	1.00
May 16, 2017 4:30:06 pm	16.00	73.00	1.00	1.00
May 16, 2017 4:29:06 pm	25.00	73.00	1.00	1.00
May 16, 2017 4:28:06 pm	13.00	73.00	1.00	1.00

Note: OmniVista can display up to 10,000 data points in a table. If OmniVista detects that the time range you are monitoring contains more than 10,000 data points, a message will display advising you to reduce the time range for the display. To change the time range, click on the time range link at the top of the table and reduce the time range. However, you can still view all of the collected data without reducing the time range by clicking on the button to view/save the table as a .csv file.

Editing a Statistics Profile

Click on a profile in the Statistics Profile List and click on the Edit icon . The Create Profile Wizard will open. Edit any parameters as described above and click on the **Update** button at the end of the wizard. You cannot edit the profile name.

Deleting a Statistics Profile

Click on a profile(s) in the Statistics Profile List and click on the Delete icon . Click on **OK** at the Confirmation Prompt. Note that deleting a profile also **deletes all statistical data** associated with the profile.

Scheduling a Profile

You can schedule a profile to start at a specific time in the future, run for a specific period of time, or run at regular intervals. If you schedule a profile to start at a specific time (set the Start Time only), the profile will start at that time and continue running. If you set a profile to run for a specific time period (set Start Time and End Time), the profile will only run during the specified time period and then stop.

If you schedule a profile to run at regular intervals, the profile will run and collect data for the profile's configured Data Retention Period, stop, and then restart at the scheduled interval. For example, if you schedule a profile with a 30-day Data Retention Period to begin running on the current date and set the Interval for 60 days, the profile will run for 30 days, and then stop. The profile will then automatically re-start after 60 days and run again for 30 days.


To schedule a profile, you must first stop the profile by selecting the profile in the Profiles List and clicking on the **Stop** icon , then click on **OK** on the Results screen. Once the profile has stopped, select it and click on the **Schedule** button. The Schedule window will appear. Configure the fields as described below to configure the profile schedule and click on **OK**.


- **Start Time** - Enter/select the date and time you want the scheduled profile to start running. If you set only a Start Time, the profile will start at the specified time and continue running. If you do not set a Start Time, the scheduled profile will begin running immediately.

- **End Time** - Enter/select the date and time you want the scheduled profile to stop running. The profile will run only for the configured time period and then stop. If you do not set an End Time, the schedule profile will continue running.
- **Interval** - Use the fields to enter a regular interval for the job to repeat (e.g., 45 days, 60 days). The profile will stop running when the Data Retention Period is reached. The profile will then automatically re-start at the specified interval. The Interval cannot be less than the Data Retention Period configured for the profile.
- **Repeat** - If you want to repeat the configured interval, enter the number of times you want it to repeat (Range = 1 - 999).

Note: You can view details of any scheduled Statistics Profile job in the User Defined Jobs tab of the Scheduler Jobs Table (Administrator - Control Panel - Scheduler Jobs).

Starting/Stopping a Profile

A profile starts collecting data as soon as it is created. You can stop a profile by selecting the profile(s) in the Profile Statistics List and clicking on the **Stop** icon  at the top of the list. Click **OK** on the Results Screen. The profile will stop collecting data. You will not lose previously-collected data. The data will be stored on the OmniVista Server based on its Data Retention Period.

To restart a profile that has been stopped, select the profile(s) and click on the **Start** icon  at the top of the list. The profile will begin collecting data and the data will be displayed in the profile. However, there will be a gap for the period during which the profile was stopped.


Statistics Profile List

- **Name** - The user-configured name of the profile.
- **Status** - The running status of the profile (Running, Stopped). If a profile is "Running", it is collecting data. If a profile is "Stopped" it is not currently collecting data.
- **Running Mode** - Indicates whether the profile is "Manual" or "Scheduled". A Manual Profile runs continuously. A Scheduled Profile runs at specified intervals.
- **Start Time** - The most recent time the profile was started and began collecting data. A profile begins collecting data when it is created.
- **End Time** - The most recent time the profile was stopped. If a profile has been continuously running since the last time it was started, the field will be blank.
- **Interval** - The profile polling rate, in seconds (Range = 1 - 60, Default = 60).
- **Data Retention** - The amount of time statistics data is saved on the OmniVista Server. (Range = 30 - 180 Days, Default = 30 Days)
- **Description** - The user-configured description for the profile.
- **Owner** - The login name of the user who created the profile (e.g., admin)
- **Created Time** - The date and time the profile was created.

Profiles

The Analytics Profiles Screen displays currently-configured Analytics Profiles, and is used to create, edit, and delete profiles. The first step in generating analytics information for Top N Applications, Top N Clients, and Top N Ports Utilization Reports is to create an Analytics Profile. A profile consists of the type of information you want to view (Profile Type) and the switches/ports that you want to analyze.

Creating a Profile


Click on the Create icon . Complete the fields in the Create Profile Wizard as described below:

Configuration Screen

- **Profile Name** - The user-configured name for the profile.
- **Profile Type** - Select a Profile Type from the drop-down menu:
 - **Top N Apps & Clients** - This profile gathers information about the top applications being accessed on the network, including which clients are accessing an application, and which switches have the most traffic for an application. Data from this profile type is displayed in both the Top N Applications Report (displays application information) and in the Top N Clients Report (displays client information).
 - **Top N Ports Utilization** - This profile gathers information about port utilization
- **Sampling Rate (Top N Apps & Clients Only)** - The ratio of packets observed at the data source to the samples generated. For example, a sampling rate of 100 specifies that, on average, 1 sample will be generated for every 100 packets observed.

Note: You can click on the **Create** button to create the profile without specifying switches/ports. At a later time you can edit the profile to add switches/ports. Otherwise, click on the **Next** button to assign the profile to the switches/ports you want to analyze.

Device/Port Selection Screen

- **Default Ports Template** - You can use this option to assign the profile to the same ports on selected switches. Rather than selecting switches and then selecting ports for each switch, you can assign the profile to the same ports on multiple switches. For example, to assign the profile to ports 1/1 through 1/10, enter 1/1-1/10, then click on the Add icon . When you select switches in the next step, the profile will be assigned to ports 1 through 10 on all selected switches and the "Apply Ports Template" message will appear.


You can also create multiple templates to be applied to different switch types. For example, you can create a template for 6850 Switches (e.g., 1/1-1/10) and create a template for 6860 Switches (e.g., 1/1/1-1/1/10). When you select switches in the next step, the profile will be assigned to ports 1 through 10 on all selected 6850 Switches, and ports 1/1/1 through 1/1/10 on all selected 6860 Switches.

- **Add/Remove Switches** - Click on the **Add/Remove Switches** button. From the list of switches, select the switch(es) you want to analyze, then click **OK**. The selected switch(es) will be displayed. Click on the **Add/Remove Ports** button to specify ports. Note that if you created a Default Ports Template, you do not need to configure ports, ports will automatically be assigned based on the template.
- **Add/Remove Ports** - Select a switch and click on the **Add/Remove Ports** button. From the list of ports, select the port(s) that you want to analyze, then click **OK**. If you selected multiple switches, select the next switch and repeat until ports have been selected for all switches. Click on the **Create** button. After clicking on the **Create** button, the status of the operation will be displayed in the Results Table. Click **OK** to return to the Profiles Screen.

Note: A switch can only be in **one** profile of a particular Profile Type.

Note: If you change the IP address of a switch after assigning a "Top N App & Clients Profile" to the switch, you must re-assign the profile to the switch.


Editing a Profile

Select the profile in the Profiles Screen and click on the Edit icon  to bring up the Edit Profile Wizard. Note that you cannot edit the Profile Name or Profile Type. Depending on the Profile Type, you can add/remove switches and/or ports to/from a profile on the Device/Port Selection Screen.

To add/remove switches to/from a profile, click on the **Add/Remove Switches** button and select the switches you want to add/remove to/from the profile. When you are done editing, click on the **Apply** button. The status of the operation will be displayed in the Results Table. Click **OK** to return to the Profiles Screen. Note that removing a switch from a profile automatically removes any ports associated with that switch and removes the sFlow configuration from the ports.

To add/remove ports from a profile, select a switch, click on the **Add/Remove Ports** button and select the ports you want to add/remove to/from the profile. When you are done editing, click on the **Apply** button. The status of the operation will be displayed in the Results Table. Click **OK** to return to the Profiles Screen. Note that removing a port from a profile removes the sFlow configuration from the port.

Deleting a Profile

Select the profile(s) in the Profiles Screen and click on the Delete icon , then click **OK** at the confirmation prompt. The status of the operation will be displayed in the Results Table. Click **OK** to return to the Profiles Screen. If the profile has been assigned to devices, a warning prompt will appear. You can click on the "Device" link to see which switches the profile has been assigned to. Click **OK** to remove the profile from the listed switches. Note that you can remove a profile from a specific switch(es) by editing the profile.

Viewing a Profile

Click on a profile on the Profiles Screen to view details of the profile. The Profile Name and Profile Type will be displayed. You can expand the Switches area to view information about switches associated with the profile.

Summary View

The Analytics Summary View Screen displays basic information for all discovered network switches, including any Analytics Profiles to which a switch might belong. Click on a switch to view detailed switch information. If the switch is included in an Analytics Profile(s), the Profile Name(s) is displayed in the Profiles field. Click on the **View** button to go to the Profiles Screen and view profile details. From the Profiles Screen, you can view, edit, or delete the profile.

Switch Information

- **Address** - The switch IP address.
- **Name** - The user-configured switch name.
- **Location** - The user-configured switch location (if no location was configured by the user, the field will display "Unknown").
- **MAC Address** - The switch MAC address.
- **Version** - The switch AOS version.

- **Type** - The switch type (e.g., OS10K, OS6900-X20).

Applications Management

When generating a Top N Applications Report, the Analytics application uses port numbers to identify application traffic. In other words, traffic on a specific port is identified as coming from a specific application. The Analytics Application Management Screen is used to create, edit, and delete application/port mapping. Well known ports (e.g., 161 for SNMP, 80 for HTTP) do not need to be mapped. By default, these ports are automatically mapped and are displayed on the screen.

Creating Application Mapping

Well known ports (e.g., 161 for SNMP, 80 for HTTP) do not need to be mapped. By default, these ports are automatically mapped and are displayed on the screen. To map other ports to an application follow the steps below.

Note: If you have an existing application ports mapping file (.json file), you can import the file rather than creating individual mappings as described in the steps below.


1. Click on the displayed mode button (Range Based/Enumerated) to select the **Mode** to use for monitoring/mapping. When you click on the button, select the **Mode** on the Select Mode Window and follow the instructions below:
 - **Range-Based** - This mode is used to set a range of ports that are monitored by the Analytics application. Traffic on these ports is monitored and can be displayed in the Top N Applications Report. Information for all of these ports is available to be displayed (depending on how you configure the report), however, only those ports that you have mapped will be labeled with the application. Other ports will be labeled "Unknown". If you select Range Based Mode, enter a range of ports to be monitored, then click **OK**.
 - **Enumerated** - This mode requires that you define specific ports to be monitored. Only those ports you define when you create a mapping will be monitored. If you select Enumerated Mode, click **OK**.
2. Click on the Create icon **+**. On the Create Application Mapping Screen, complete the fields as described below:
 - **Application Name** - Enter the name of the application (e.g., SNMP) .
 - **Ports** - Enter the port or port range to be associated with the application. If you are entering a range of ports, separate the port numbers with a "-" (e.g., 20-21).

Importing/Exporting an Application Ports Mapping File


If you have an existing application ports mapping file (.json file), you can import the file into OmniVista 2500 NMS. Click on the **Import** button to bring up the Import an Applications Ports Mapping File window and click on the **Browse** button. Locate the file and click **OK**. The port mappings in the file will appear in the list on the Applications Management Screen. Note that this new mapping **will override** your existing mapping.

You can also create an application ports mapping .json file by exporting your existing mapping list. To create/export the file, click on the **Export** button. At the prompt screen, select **Save File** and click **OK**. The file will be downloaded to your default download area.

Editing Application Mapping


Click in the checkbox next to a mapping entry and click on the Edit icon  to bring up the Edit Application Mapping Screen. Edit the Application Name and click on the **Update** button. The updated entry will appear in the list on the Application Management Screen. You cannot edit the ports. To map a different Application to a port, you must delete the mapping entry and create a new one.

Deleting Application Mapping

Click in the checkbox next to a mapping entry and click on the Delete icon , then click **OK** at the confirmation prompt.

Anomalies

The Analytics Anomalies Screen displays any anomalies that are discovered in established port utilization trends. The information is displayed in a list that describes the anomaly and its origins (e.g., IP address, Port). Anomaly detection uses Z-Score to check for anomalies in the latest port utilization data gathered from hourly polling over the past 30 days. Z-Score is a statistical measurement of a score's relationship to the mean in a group of scores. In other words, it measures utilization for a port for a specific hour to determine its relationship with utilization for the same hour over the sampling period (30 days). A data point that deviates considerably from an established pattern is flagged as an anomaly and displayed on the Anomalies Screen. Z-Score parameters are configured on the Preferences - Analytics Screen.

You can configure the information displayed by clicking on the Configuration icon  to bring up the Configuration Screen and set any or all of the displayed columns. Click on the **Add To Report** button to create a report in the Report Application (see the Report Configuration Help for more information).

Note: A minimum of 11 days of data is required for anomaly calculation. Also, seasonal variation for periods of more than 30 days cannot be adequately learned using this method. For example, an annual usage pattern would be affected by lower usage due to holidays/vacations.

Settings

The Analytics Settings Screen is used to configure preferences for port utilization trending and anomaly detection in the Analytics application, as well as preferences for Analytics Statistics. When you have configured the value(s), click the **Apply** button. The change takes effect immediately.

Analytics Configuration

- **sFlow Port** - The sFlow port used to gather analytics data (Default = 6343).
- **Outlier Detection: Lower Threshold** - Used for anomaly detection. The threshold value used to determine if new port utilization is lower than mean utilization for that port:
 - BEYOND2Z-SCORE - Any value that falls outside 2 times Standard Deviation from mean on a normal distribution curve.
 - BEYOND2.5Z-SCORE - Any value that falls outside 2.5 times Standard Deviation from mean on a normal distribution curve.
 - BEYOND3Z-SCORE - Any value that falls outside 3 times Standard Deviation from mean on a normal distribution curve. (Default)
- **Outlier Detection: Higher Threshold** - Used for anomaly detection. The threshold value used to determine if new port utilization is higher than mean utilization for that port:
 - BEYOND2Z-SCORE - Any value that falls outside 2 times Standard Deviation from mean on a normal distribution curve.

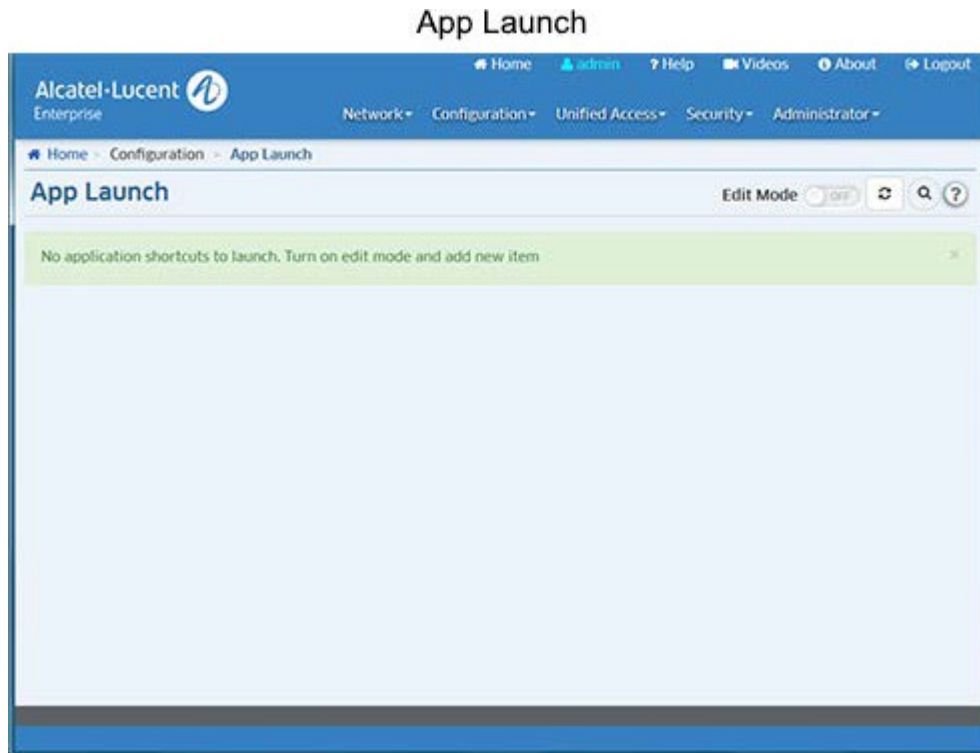
- **BEYOND2.5Z-SCORE** - Any value that falls outside 2.5 times Standard Deviation from mean on a normal distribution curve.
- **BEYOND3Z-SCORE** - Any value that falls outside 3 times Standard Deviation from mean on a normal distribution curve. (Default)
- **Prediction: Training Timeout** - Used for port utilization trending. Specifies how long OmniVista will train, in seconds, by sampling past port utilization. In other words, this specifies how long OmniVista will sample port utilization data before beginning to predict future trends (Range = 15 - 600, Default = 60).
- **Prediction: Training Error** - Used for port utilization trending. The target error percentage to which OmniVista will be trained (Default = 0.1 - 1.0, Default = 0.5).
- **Top N Ports Purge** - The amount of time, in months to retain analytics port utilization data before it is purged from the OmniVista Database (Range = 1 - 8, Default = 6).
- **Top N Switches Purge** - The amount of time, in months to retain analytics switch resource data before it is purged from the OmniVista Database (Range = 1 - 8, Default = 6).
- **Top N Apps Purge** - The amount of time, in months to retain analytics application data before it is purged from the OmniVista Database (Range = 1 - 100, Default = 6).
- **Top N Clients Purge** - The amount of time, in months to retain analytics client data before it is purged from the OmniVista Database (Range = 1 - 100, Default = 6).

Application Visibility Statistics Configuration

- **AppMon/AppFP Schedule Interval** - The amount of time OmniVista will wait to poll devices for Application Monitoring Statistics.
- **DPI Schedule Interval** - The amount of time OmniVista will wait to poll devices for Application Enforcement Statistics.

4.0 App Launch

The App Launch Screen enables you to launch web-based (e.g., OpenStack) applications using OmniVista. You can add/edit/delete application links and arrange the links on the page. Once a link is added to the page, you can click on it to open the application in a new browser tab.



Edit Mode


To add/edit/delete an application link or arrange links on the page, click on the **Edit Mode** slider to change the mode to **On**. The Create icon **+** will appear next to the slider, and any existing links on the page will display in Edit Mode.

Adding a Launch Icon


To add a launch icon, click on the Create icon **+** to bring up the "Add New Application" Window. Enter an **Application Name** and the **URL** needed to access the application. If you have an image for the icon, click on the **Browse** button to locate the image file. (If you do not have an image, a generic image will be used along with the Application Name you entered.) When you are done, click on the **Add** button. The icon will appear on the App Launch page. When you are finished, change the **Edit Mode** slider back to **Off**.

Note: Images can be .jpg, .gif, or .png files, with a maximum size of 60 x 60 pixels. Note that if you add a large number of links you can use the search feature to search for a specific application link on the screen. Enter the name of the application in the **Search** field. The other links are temporarily removed and the link you are searching for is isolated on the screen.


Editing a Launch Icon

To edit a launch icon, change the **Edit Mode** to **On** and click on the edit symbol  in the corner of the icon to bring up the "Edit Application" Window. Edit the necessary fields and click on the **OK** button. The updated icon will appear on the App Launch page. When you are finished, change the **Edit Mode** to **Off**.

Deleting a Launch Icon

To delete a launch icon, change the **Edit Mode** to **On** and click on the delete symbol  in the corner of the icon. Click **Yes** at the confirmation prompt. When you are finished, change the **Edit Mode** to **Off**.

Arranging Launch Icons

Change the **Edit Mode** to **On**, you can also arrange the icons on the page. Click on the **Setting** icon  to bring up the "Settings" Screen. Configure the display options as described below.

- **Items per Page** - Select the maximum number of icons you want displayed on the page (15, 30, 45).
- **Sort By** - Select **Name** to display the icons alphabetically by name. Select **Creation Date** to display the icons chronologically by when they were added to the page. Select **Custom** to sort the icons manually by dragging them to new positions.

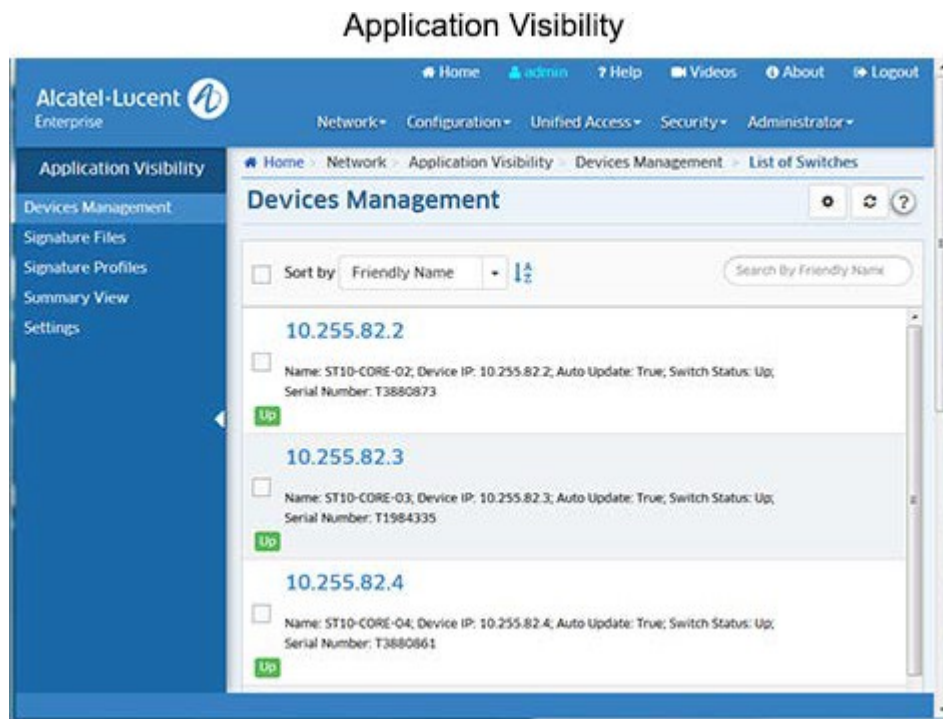
After making your selections, click **OK**. If you selected **Name** or **Creation Date**, the icons will be automatically displayed in the order selected. If you selected **Custom**, you can then drag the icons to their new positions. When you are finished, change the **Edit Mode** to **Off**.

5.0 Application Visibility

The Application Visibility Application supports monitoring and QoS configuration of Application traffic flows, and performs statistics profiling on the collected data. OmniVista simplifies Application Visibility configuration for the network by enabling you to quickly configure Application Visibility on switches throughout the network.

Application Visibility is supported on OS10K Switches (AOS 7.3.4.R01 and later), OS6900 Switches (AOS 7.3.4.R02 and later), and OS6860E Switches (AOS 8.2.1.R01 and later). It is also supported in a virtual chassis of OS6860/OS6860E Switches where at least one OS6860E is present.

Note: OmniVista must be able to FTP to a switch to gather information for the Application Visibility Application. When you initially discover network switches using the Discovery Application, make sure to specify the CLI/FTP User Name and Password for the switches in the Discovery Profile. If you do not specify the CLI/FTP User Name and Password during discovery, you can specify them anytime using the Topology Application. Go to the Topology application (Network - Topology), select the switch in the Topology Map that you want to monitor, click on the **Discovery - Edit Device** option in the Operations panel to bring up the Edit Discovery Manager Device Screen and edit the fields.



Application Visibility Application

The Application Visibility Application identifies application/protocol flows based on Application Signatures that identify an associated application or protocol. To enable Application Visibility, you create a Signature Profile that includes all of the signatures for applications/protocols that you want to monitor/control, and apply that profile to network switches/ports. Once Application Visibility is enabled on switch ports, the switch can identify traffic flows included in the profile. You can then use the Top N Applications - Advanced Report in the Analytics application to monitor network usage for each application/protocol; and use the Application Enforcement feature in Application Visibility to assign QoS/UNP Policies to the traffic.

The page links on the left side of the screen (shown above) are used to view and configure Application Visibility:

- **Devices Management** - Displays all network switches supporting Application Visibility. In addition to the name, IP address, and operational status of each switch, the screen indicates whether or not an Application Visibility Profile has been assigned to the switch. You can also select a switch to display more detailed information and enable/disable automatic Signature Profile updates. (Automatic Signature Profile updates are only supported for 8.x Signature Files (OS6860/6860E Switches).
- **Signature Files** - Displays all Signature Files downloaded/imported into OmniVista. It is also used to download/import Signature Files.
- **Signature Profiles** - Used to create Signature Profiles. Signature Profiles are created from a Signature File, which contains Application Signature information in pre-configured Application Groups (groups of related applications protocols). You create a Signature Profile by selecting one or more applications/application groups (or creating a custom group) that contain the applications/protocols you want to monitor/control. You then assign the Signature Profile to network switches/ports.
- **Summary View** - Used to view information on switches configured for Application Visibility.
- **Settings** - Used to configure Signature File update settings. Signature Files are regularly updated to either provide new signatures, or to update existing signatures which have changed. The Settings Screen is used to configure how often OmniVista will check the ALE Signature File Repository for updates. If an update is available, OmniVista will automatically download the file. You can also configure how often OmniVista will check any switches configured for Application Visibility to make sure they have the most recently downloaded Signature File. Automatic Signature File updates are only supported for 8.x Signature Files (OS6860/6860E Switches).

Application Visibility Configuration

The Application Visibility Application supports monitoring and QoS/UNP configuration of Application traffic flows, and performs statistics profiling on the collected data. OmniVista simplifies Application Visibility configuration for the network by enabling you to quickly configure Application Visibility on switches throughout the network. Configuring Application Visibility for Application Monitoring and Application Enforcement consists of the following steps:

1. Download and import Signature Files from the Alcatel-Lucent Enterprise (ALE) Signature File website. There are different Signature Files for each of the supported device types. Signature Files for OS10K/6900 Switches are developed/updated by the user and can be imported into OmniVista. Signature Files for OS6860/6860E Switches are provided by ALE and are automatically downloaded/updated. (Signature Files Management Screen)

2. Create Signature Profiles from the Signature File. (Signature Profiles Screen)

- **OS10K/6900 Switches**
 - **Configure Monitoring** - Using the "Create Signature Profile" Wizard. Wizard guides the user through configuration of application monitoring groups.
 - **Configure Enforcement** - Using the "Configure UNP/QoS" Wizard. The wizard provides a link to the PolicyView Application for creation of Application Visibility Policies/Policy Lists that are used for QoS enforcement. The wizard also provides a link to the Unified Access Application for creation of Access Role Profiles for UNP enforcement.

- **OS6860/6860E Switches**

- **Configure Monitoring** - Using the "Create Signature Profile" Wizard. Wizard guides the user through configuration of application monitoring groups.
- **Configure Enforcement** - Using the "Create Signature Profile" Wizard. Wizard guides the user through configuration of application enforcement groups. User must then configure an Access Role Profile in the Unified Access Application based on the configured enforcement groups.

3. Apply Signature Profiles to network switches/ports. (Signature Profiles Screen)

Devices Management

The Application Visibility Devices Management Screen displays all network switches that support Application Visibility. In addition to the name, IP address, and operational status of each switch, the screen indicates whether or not an Application Visibility Profile has been assigned to the switch. Click on a switch to display more detailed information and/or enable/disable automatic Signature File updates. Note that the Signature File Auto-Update Feature is supported on OS6860/6860E Switches only.

Device Information

You can view detailed information for Switches by clicking on a switch.

- **Name** - The switch name, if applicable.
- **Device Version** - The switch AOS build number.
- **Serial Number** - The switch serial number.
- **Type** - The switch model type (e.g., OS6900-X72).
- **Device IP** - The device IP address ("Master" switch/chassis IP address for a stack or virtual chassis configuration).
- **MAC Address** - The device MAC address ("Master" switch/chassis MAC address for a stack or virtual chassis configuration).
- **Signature File** - The name of the Signature File assigned to the switch, if applicable.
- **Signature Version** - The version of the Signature File assigned to the switch, if applicable.
- **Signature Profile** - The name of the Signature Profile assigned to the switch, if applicable.
- **Switch Status** - The switch operational status.
- **Sync Status** - Whether or not the Signature Profile assigned to the switch is in sync with the profile stored in OmniVista. If the Signature Profile in OmniVista is different than the one on the switch, Sync Status will indicate "out of sync".
- **Auto Update (OS6860/6860E Only)** - Enables/Disables (On/Off) Signature File auto update. If the "Signature Auto Update" option is enabled on the Application Visibility Settings Screen, OmniVista automatically downloads newer 8.x Signature File versions from the ALE Signature File Repository when they become available. If Signature File "Auto Update" is enabled on a switch, after downloading a new Signature File version, OmniVista first updates any Signature Profiles created with the older Signature File version. If that updated Signature Profile is being used by the switch, OmniVista then automatically assigns the updated profile to switch.

Note that auto update only occurs on Signature Files within the same "Major" version (the first number in the file version). For example, if version 1.1.1 is stored in OmniVista and version 1.2.1 is available in the repository and downloaded, OmniVista will automatically update the Signature Profile(s) that are using Signature File version 1.1.1 with version 1.2.1 and apply them to any switches with that profile. If version 2.1.1 is available and downloaded, OmniVista will download the file, but will not update any profiles using 1.x.x Signature Files switches using

those profiles.

Signature File automatic update check frequency is configured on the Application Visibility Settings Screen.

Signature Files

The Application Visibility Signature Files Screen displays all imported Signature Files, and is used to upgrade/import Signature Files. A Signature File contains application signature information that is used to create Signature Profiles. Once you create a Signature Profile, you assign that profile to switches to monitor/control application traffic on the network.

Viewing Signature Files

The Signature File Management Screen displays all Signature Files in OmniVista. There are different Signature Files for each of the supported device types. Signature Files for OS10K/6900 Switches are .txt files (e.g., 7xSignatureFile.txt) that are developed/updated by the user and can be imported into OmniVista. Signature Files for OS6860/6860E Switches are binary files (e.g., UAppSig.upgrade_kit_1) that are provided by Alcatel-Lucent Enterprise (ALE) and are automatically downloaded/updated.

Note: When you initially install OmniVista, you must create and import 7.x Signature Files into OmniVista for OS10K/6900 Switches. For 8.x Signature Files (OS6860/6860E Switches), you must go to the Application Visibility Settings Screen and click on the **Update Now** button to upload the initial Signature File to OmniVista. After this initial upload, if the **Signature Auto Update** option is **Enabled** on the Application Visibility Settings Screen, OmniVista automatically downloads newer 8.x Signature File versions when they become available. Note that the "Signature Auto Update" option is disabled by default. You must enable it for automatic updates.

Importing a Signature File

OmniVista automatically checks the ALE Signature File Repository and updates and downloads 8.x Signature Files for OS6860/6860E Switches (see Upgrading a Signature File below). There should be no need to import these Signature Files into OmniVista. If necessary, you can perform a manual upload by going to the Application Visibility Settings Screen and clicking on the **Update Now** button.

Signature Files for OS10K/6900 Switches are .txt files (e.g., 7xSignatureFile.txt) that are developed/updated by the user and must be imported into OmniVista. After creating/updating the file, download the file to your computer, then click on the **Import File** button to import the file into OmniVista. The Import Signature File Screen will appear. Click on the **Browse** button to locate the file(s) and click **OK**.

Upgrading a Signature File


OmniVista regularly checks the ALE Signature File Repository for newer Signature File versions (the update check frequency is configured on the Application Visibility Settings Screen.) If OmniVista detects that a new Signature File version is available (e.g., version 1.1.1 is stored in OmniVista and version 1.2.1 is available in the repository), OmniVista automatically downloads the file. However, you can manually update a Signature File at any time by selecting the file and clicking on the **Upgrade** button.

Note: OS6860/6860E Switches support the Signature File Auto-Update Feature. If Signature File "Auto Update" is enabled on a switch (configured on the Devices Management Screen), after downloading the new Signature File version, OmniVista first updates any Signature Profiles created with the older Signature File. If a switch has "Auto Update" enabled and that

Signature Profile is being used by the switch, OmniVista then automatically assigns the updated profile to switch. Note that profiles and switches are only automatically updated if they are using the same "Major" version of the new Signature File (e.g., profile and switch are using version 1.1.1 and version 1.2.1 is available and downloaded by OmniVista).

For OS10K/6900 Switches, you must manually update the affected Signature Profile and re-assign the profile to the switch.

Deleting a Signature File

To delete an imported Signature File from the repository, select the file in the table and click on the Delete icon . Click **OK** at the confirmation prompt. Note that you cannot delete a Signature File that has been assigned to switches on the network.

Signature Profiles

The Application Visibility Signature Profiles Screen displays all configured Signature Profiles and is used to create, apply, edit, clone, and delete profiles. Signature Profiles are created from a Signature File, which contains Application Signature information for individual applications/protocols as well application groups (pre-configured groups of related applications/protocols). You create a Signature Profile by selecting one or more applications/application groups (or creating a custom group) that contain the applications/protocols you want to monitor/control. You then assign the Signature Profile to network switches. Multiple Signature Profiles can be created from a single Signature File, each containing a different combination applications/application groups. And a Signature Profile can be assigned to one or more switches. However, a switch can be assigned only one Signature Profile.

Note: Application/protocol traffic is monitored using the Analytics Application. To view statistics on applications/protocols you have configured in a profile, go to the "Top N Applications - Advanced" Screen (Network - Analytics - Top N Applications - Advanced). Once you configure a profile and assign it to switches statistics for the applications/protocols in the profile are displayed in graphical and table format.

Viewing Signature Profiles

The Signature Profiles Screen displays all configured Signature Profiles. Click on a profile to display detailed profile information.

- **Profile Name** - The user-configured name for the profile.
- **Description** - A user-configured description for the profile.
- **File Name** - The name of the Signature File used in the profile.
- **File Version** - The file version of the Signature file used in the profile.
- **Apps** - Lists the applications included in the profile.
- **Groups** - Lists the application groups included in the profile.
- **Devices** - Lists the switches to which the profile has been assigned.

Creating a Signature Profile

Signature Profiles are created from a Signature File, which contains application signature information for individual applications/protocols and application groups. You create a Signature Profile by selecting one or more applications/application groups (or creating a custom group) that contain the applications/protocols you want to monitor/control. The Signature Profile Wizard guides you through the steps to create a profile. Because the monitoring and enforcement implementation differs depending on switch type (7.x - OS10K/6900 vs. 8.x - OS6860/6860E), the steps within the wizard are slightly different depending on

which Signature File Type you use.

For OS10K/OS6900 Switches, you configure application monitoring groups using the "Create Signature Profile" Wizard. Application enforcement is configured using the "Configure UNP/QoS Wizard", which you access by clicking on the **Configure UNP/QoS** button at the top of the Signature Profiles Screen. For QoS, the wizard guides you through the process of configuring Application Visibility Policy Lists using Signature File Groups. You then apply the Application Visibility Policy List to network switches/ports. For UNP, the wizard is used to enable Application Visibility UNP on switch ports. UNP is configured separately using CLI Scripting.

For OS6860/OS6860E Switches, you configure monitoring and enforcement for individual applications or application groups using the "Create Signature Profile" Wizard. For QoS/UNP, you then create an Application Visibility Policy List using Signature File groups; and create an Access Role Profile using the Policy List.

The sections below detail configuring a profile for OS10K/6900 Switches and OS6860/6860E Switches.

OS10K/6900 Switches

Follow the steps below to create a Signature Profile for OS10K/6900 Switches. There must be at least one application group in a profile. For OS10K/OS6900 Switches, the groups you use when configuring a profile are used for both monitoring and enforcement. The "Create Signature Profile" Wizard is used to configure application monitoring. The "Configure UNP/QoS" Wizard is used to configure application enforcement.

Configuring Application Monitoring

The "Create Signature Profile" Wizard for OS10K/OS6900 Switches guides you through creating a Signature Profile. Click on the Create icon **+** to create a new profile. The "Create Signature Profile Wizard" appears.

Complete the screens as described below. After creating a profile, you must apply it to switches/ports in the network.

Name and Description

Enter a **Profile Name**. You will be using different Signature Files to create profiles for OS10K/OS6900 Switches and OS6860/OS6860E Switches, so you should enter a name describing which profile type you are creating (e.g., 7.x Profile, 8.x Profile). You can also enter a profile **Description**. Click **Next**.

Select File

Select a 7.x Signature File (e.g., 7xSignatureFile.txt), then click **Next**.

Select Groups

Click on the **Choose App Groups** button to select the groups you want to include in the profile, click **OK**, then click on the **Create Profile** button. You can also create a custom Application Group to include only those applications that you want to monitor by clicking on the **Create App Group** button. Enter an **Application Group Name** and **Description**, select the applications you want to include in the group, click **OK**, then click on the **Create Profile** button.

After creating a profile, you must apply it to switches/ports in the network. Configuration is now complete for Application Monitoring. You can view application flow information in the Analytics Application. At this point, you can also configure application enforcement by clicking on the **Configure UNP/QoS** button at the top of the screen to bring up the "Configure UNP/QoS" Wizard.

Configuring Application Enforcement

The "Configure UNP/QoS" Wizard for OS10K/OS6900 Switches guides you through configuring application enforcement using the PolicyView and Unified Access Applications. Click on the **Configure UNP/QoS** button at the top of the screen to bring up the wizard, and complete the screens as described below.

Select File


Select a 7.x Signature File (e.g., 7xSignatureFile.txt), then click **Next**.

Select Signature Profile

Select the 7.x Signature Profile you want to use, then click **Next**.

Configure QoS

If you have already configured Application Visibility Policy Lists, select the list you want to use, then click on the **Add/Remove Ports** button to bring up the Ports Selection Window. The window displays all of the ports for the switches contained in the Policy List. Select the ports to which you want to apply the policy list and click **OK**. Click the **Apply** button to apply the QoS configuration, or click the **Next** button to configure UNP (optional).


Note: If you have not yet configured Application Visibility Policy Lists, click on the **Go to Policy List** button to go to the Application Visibility Policy Lists Screen and configure an Application Visibility Policy List. Click on the Add icon  next to the Add Application Visibility Policies field to create an Application Visibility Policy(ies) for the policy list. The policies must be created for the same switches to which you applied the Signature Profile. (On the **Device Selection** Screen, select the same switch(es) to which you applied the Signature Profile.) On the **Set Condition** Screen, select an application group.

Configure UNP

If you have already configured Unified Network Profiles (UNP), select a switch and click and click on **Add/Remove Ports** button to bring up the Port Selection Window for the switch. Use the **Add/Remove** buttons to select the ports on which you want to enable the UNP. Click **OK**. If necessary, repeat for additional switches, then click the **Apply** button. The status is displayed on the Configure UNP/QoS Screen. Click **OK** to return to the Signature Profiles Screen.

Note: UNP is configured using the CLI Scripting application. Configuring UNP is optional. You can configure and apply QoS without configuring UNP. If necessary, go to the CLI Scripting application to configure UNP.

OS6860/6860E Switches

Follow the steps below to create a Signature Profile for OS6860/6860E Switches. There must be at least one application/application group in a profile. In addition to monitoring groups, you can also create enforcement groups for OS6860/6860E Switches using the "Create Signature Profile" Wizard. Click on the Create icon  to create a new profile. The "Create Signature Profile" Wizard appears. Complete the screens as described below. After creating a profile, you must apply it to switches/ports in the network.

For OS6860/OS6860E Switches, the wizard guides you through creating a Signature Profile containing both monitoring groups and enforcement groups. You can create monitoring groups only, without creating enforcement groups. However to configure enforcement, you must configure an enforcement group in the wizard. For enforcement, you then create an Application Visibility Policy List that you use to configure an Access Role Profile.

Name and Description

Enter a **Profile Name**. You will be using different Signature Files to create profiles for OS10K/OS6900 Switches and OS6860/OS6860E Switches, so you should enter a name describing which profile type you are creating (e.g., 8.x Profile). You can also enter a profile **Description**. Click **Next**.

Select File

Select an 8.x Signature File (e.g., UAppSig.upgrade_kit_1), then click **Next**.

Select Monitoring Groups and Apps

Click on **Groups**, then click on the **Choose App Groups** button and use the **Add/Remove** buttons to select the groups you want to include in the Monitoring Profile, and click **OK**. You can also create a custom Application Group to include only those applications that you want to monitor by clicking on the **Create App Group** button. Enter an **Application Group Name** and **Description**, use the **Add/Remove** buttons to select the applications you want to include in the group, and click **OK**. Note that will be helpful to add a descriptive name to the custom group to easily identify in it case you want to configure application enforcement (e.g., 8,xEnforcement).

On OS6860/6860E Switches you can also configure individual applications in the profile. Click on **Apps**, then use the **Add/Remove** buttons to select the applications you want to include in the Monitoring Profile and click **OK**. Note that if an application is included in a group, you cannot configure it individually.

At this point, you can click on the **Create Profile** button to just create a Monitoring Profile, or click the **Next** button to configure an Enforcement Profile.

Select Enforcement Groups and Apps


Click **Groups**, then click on the **Choose App Groups** button and use the **Add/Remove** buttons to select the groups you want to include in the Enforcement Profile, click **OK**, then click the **Create Profile** button. You can also create a custom Application Group to include only those applications that you want to control by clicking on the **Create App Group** button. Enter an **Application Group Name** and **Description**, use the **Add/Remove** buttons to select the applications you want to include in the group, click **OK**, then click on the **Create Profile** button.

The profile creation is complete and can be used to create reports for monitoring the applications in the profile using the Analytics Application. To configure application enforcement, and create Application Count Reports in the Analytics application, you must create an Application Visibility Policy List using Signature File Groups; and create an Access Role Profile using the Policy List.


1. Create an Application Visibility Policy(ies) for the same switches to which you applied the Signature Profile, using an Enforcement Group/App contained in the Signature Profile. (Configuration - PolicyView - Application Visibility Policy)


- On the **Device Selection** Screen, select the same switch(es) to which you applied the Signature Profile. On the **Set Condition** Screen, select one of the groups or applications included in the Signature Profile. As mentioned earlier, when creating an Enforcement Group, it is recommended that you enter a name describing the group so it is easily identified here (e.g., 8.xEnforcement).
2. Create a Policy List containing the Application Visibility Policy(ies) you created in Step 1. (Configuration - PolicyView - Policies & Policy Lists - Policy List).
 3. Create an Access Role Profile using the Policy List you created. (Unified Access - Unified Profile - Access Role Profile)
 - In the **Policy List** field, select the Policy List you created in Step 2. (Unified Access - Unified Profile - Access Role Profile).
 4. Apply the Access Role Profile to the same switches to which you applied the Signature Profile.

Editing a Signature Profile

From the Signature Files Management Screen, click on an Upgrade Kit to display the Signature Files. Select a Signature File and click on the Edit icon . You can edit the Profile Name, Description, and Application Groups as described above. When you are done editing, click on the **OK** button. After editing the profile, you must apply it to switches/ports in the network. Note that you cannot edit a Default Profile.

Cloning a Signature Profile

You can clone an existing profile and edit it to create a new profile. Note that when you import a Signature File, a Default Profile is created for that file type (e.g., 7.x or 8.x) and appears in the Signature Profiles Table. These profiles contain all of the applications/application groups for each file type. You can create a new profile from scratch as described below, or you can clone one of the default profiles and modify it to create a new profile. To clone and modify a Default Profile, select the profile and click on the Clone icon . The "Create Signature Profile Wizard" appears. Use the wizard to modify the default profile to create a new one.

1. Select a Signature File and click the Clone icon .
2. Edit the profile as described above and click the **OK** button.
3. Apply the profile to switches/ports in the network.

Applying a Signature Profile

After creating/editing/cloning a profile, you must apply it to switches/ports in the network. Select the Signature Profile and click on the **Apply** button at the top of the screen. The Apply to Devices Screen will appear.

1. Click on the **Add/Remove Switches** button to bring up the Device Selection window. The available switches will be displayed on left side of the screen. Only switches without an applied Signature Profile that support the profile type you are applying are displayed (e.g., only 8.x switches are displayed for a profile created with an 8.x Signature File). If the profile has already been applied to a switch, the switch will already appear on the right.
2. Use the **Add/Remove** buttons to select the switches to which you want to apply the profile, then click **OK**. You will be returned to the Device Selection window.

3. Select a switch and click on the **Add Port** link or the **Add/Remove Monitoring Ports** button to bring up the Ports Selection window. 7.x Switches support Application Visibility configuration on link aggregates; 8.x switches do not.

4. Use the **Add/Remove** buttons to select the ports to which you want to apply the profile, then click on the **Apply** button. The progress is displayed on the Action Results Screen. Click **OK** to return to the Signature Profiles Screen.


Note: You can only assign one (1) Signature Profile to a switch. Also, when you apply a Signature Profile, any pre-existing Application Visibility configuration on a switch is erased and the new profile configuration is used, including any Application Visibility configuration done from the CLI.

Note: To apply a new profile to a switch with an existing profile, you must first remove the old profile from the switch before assigning the new one.

Removing a Signature Profile

As mentioned above, when applying a profile only supported switches **without** an assigned Signature Profile are displayed. If you want to apply a different profile to a switch, you must first remove the old profile to apply a new one. The process is similar to applying a profile. Select the Signature Profile you want to remove from the switch and click on the **Apply** button. Click on the **Add/Remove Switches** button to bring up the Devices Selection window. The switches to which the profile has been applied appear on the right side of the window. Use the **Add/Remove** buttons move the switch to the left side of the screen, then click **OK**. Click on the **Apply** button.

Deleting a Signature Profile

To delete a profile, select it and click on the Delete icon . Click **OK** at the confirmation prompt. Note that you cannot delete a Signature Profile that has been applied to devices on the network. You must first remove the profile from any devices before deleting it.

Summary View

The Application Visibility Summary View Screen is used to view information on switches configured for Application Visibility. Click on a switch to view which ports are enabled for monitoring/enforcement.

- **Friendly Name** - The device IP address.
- **Name** - The user-configured switch name.
- **MAC Address** - The switch MAC address.
- **Version** - The AOS software version installed on the switch (e.g. 8.2.1.309.R01).
- **Location** - The physical location of the switch (e.g., Lab).
- **Status** - The administrative status of the switch (Up/Down).
- **Type** - The switch model type (e.g., OS6860E-U28).
- **DNS Name** - The switch DNS name.
- **File Name** - The name of the Signature File contained in the Signature Profile (e.g., UAppSig.upgrade_kit)
- **File Version** - The Signature File version (e.g., 1.1.2).
- **Profile Name** - The name of the Signature Profile assigned to the switch.

Settings

The Application Visibility Settings Screen is used to configure automatic Signature File update settings. Signature Files are regularly updated to either provide new signatures, or to update existing signatures which have changed. OmniVista can be configured to periodically check the ALE Signature File Repository to determine if a new Signature File is available. If "Signature Auto Update" is enabled, OmniVista will check the ALE Signature File Repository as configured below. If a new file version is available, OmniVista will automatically download the file and update any Signature Profiles using the older Signature File version. If "Auto Update" is enabled on a switch (configured on the Devices Management Screen), OmniVista will automatically update the switch if it is using the updated Signature Profile.

You can also check for a recent Signature File update any time by clicking on the **Update Now** button. And you can click on the **Test Connection** button to check if you have properly configured the URL Signature File Repository connection.

Configure the fields as described below. When you are finished, click on the **Save** button to save the new preferences to the OmniVista Server. Click on **Revert** to return a field to its previous value. Click on **Default** to return all values to the default settings.

Note: Automatic Signature File updates are only supported on 8.x Signature Files (OS6860/6860E Switches).

Audit Configuration

- **Audit Switch Every** - How often OmniVista will check the Signature File version on Application Visibility-configured switches, in Hours. The audit is used to verify that Signature Files contained in a Signature Profile on a switch are in sync with the Signature Files stored in OmniVista. If the profile on a switch is out of sync with the Signature File in OmniVista, the status for the switch will be changed to "Out of Sync". (Range = 1 - 24, Default = 1)

Update Configuration

- **Check for Every** - How often OmniVista will check the Signature File Repository for updates (1 day, 7 days, 15 days, 30 days, None). If an update is available, OmniVista will automatically download the Signature Files. OmniVista will download the latest Signature Files from the Repository when it runs the first time.
- **Update Time** - The number of hours after a new Signature File is downloaded that OmniVista will wait before updating the file on applicable Signature Profiles/switches.
- **Signature Repository** - The location of the Signature File Repository (pre-filled - <https://ep1.fluentnetworking.com/omnivista/signature/pull>).
- **User Name** - The username for the Signature File Repository (pre-filled - *omnivista*).
- **Password** - The password for the Signature File Repository (pre-filled).
- **Signature Auto Update** - Enables/Disables (On/Off) automatic Signature File check/download.

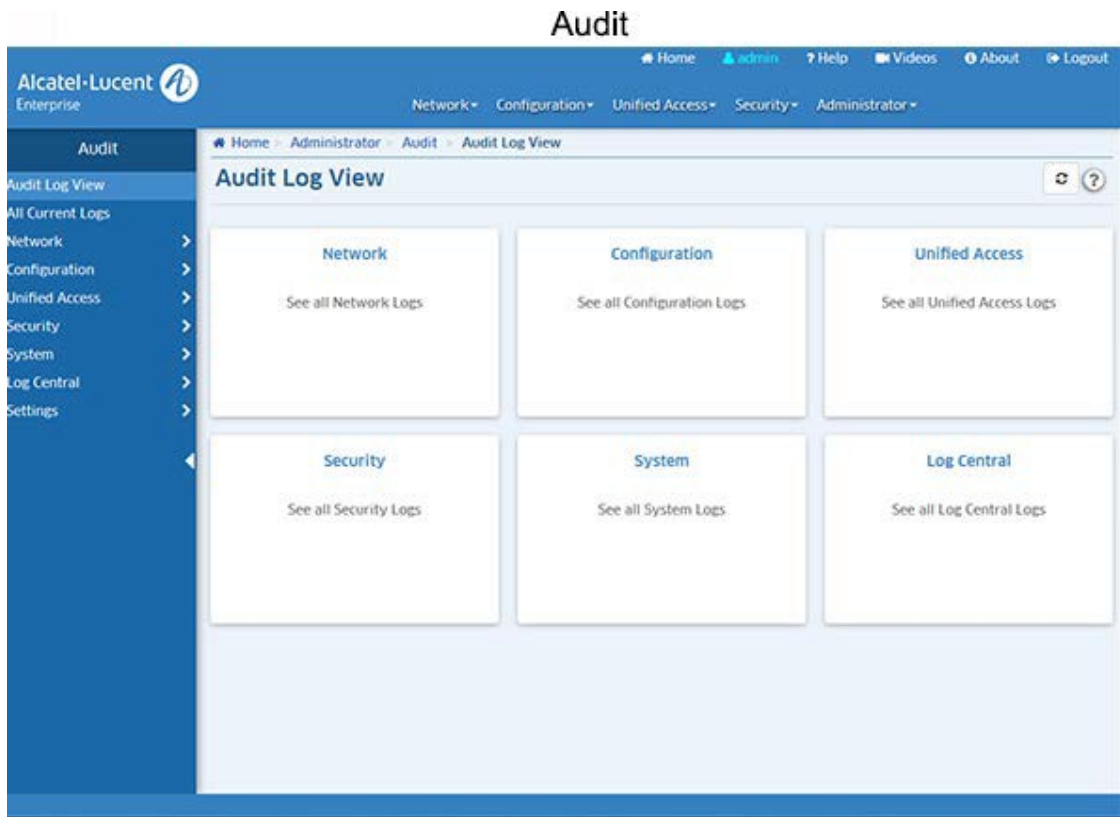
Notes:

- If OmniVista fails to download signature files, it will retry 5 times. The interval between the retries is 5 minutes.
- OmniVista will log all actions (e.g., check update, download file, update to switches) in the "afn_autoupdate.log" file.

6.0 Audit

The Audit application is used to monitor client and server activity, such as the date and time when a user logged into OmniVista, when an item was added to the discovery database, when a configuration file was saved, or when a particular application was launched. The information is contained in log files, which are organized by type (e.g., Network, Configuration), as shown on the Audit Home Page below.

The application enables you to view log files, search through log files, and download files. You can view current log files or historical log files that have been archived by OmniVista. Log Files are archived based on preferences configured on the Audit Settings Screen.




Log Files by Type

You can click on one of the tiles on the home page to view the log files in a category (e.g., Network, Configuration), or you can click on a category on the left side of the screen to display a list of log files in the category. Click on a log file to display the contents of the file. The log files in each category are listed below.

Note: Click on the "All Current Logs" link on the left side of the home page to display a list of all current logs. You can then click on a log to display the contents.

Network	Configuration	Unified Access	Security	System
Analytics Predictive Service	CLI	Access Guardian 2.0	Quarantine	Active MQ
Analytics Service	Scripting	BYOD	SecureView SA	DAL Service
Analytics Service Worker	Policy	mDNS	Security	FTP Service
AV Audit Service	Resource Manager	Wireless Service	Configuration	HSQL DB
AV Auto Update	Resource Manager			Master Poller Service
AV License	Backup Info			Mongo DB
AV Service	Resource Manager			Open LDAP
Call Home Service	Client			OV 2500 Server
Discovery	Service SIP			OV Client
Discovery Lite	Service			Redis
Polling	VLAN Creation Result			Scheduler Service
Trap Config	VLAN Deletion Result			Syslog
Traps	VLAN Service			Tomcat Catalina
VM Locator	VLAN Service Worker			Tomcat Host Manager
VM Manager Error VM Manager Service	VXLAN Service			Tomcat Local Host
				Tomcat Local Host Access Log
				Tomcat Manager
				Tomcat OV Report
				Tomcat OV Web
				Watchdog CLI
				Watchdog Service
				Worker Poller Service

Note: The Log Central link displays the ngnms.log file. This file includes log entries from all log files in real time, with the most recent entries at the top of the file. You can scroll through the file using the scroll bar and arrows on the right side of the screen, or you can also search by keyword. Enter the search criteria in the "Search in this page" field at the top of the log file and click on the **Search Next** button. The word or phrase is highlighted in yellow throughout the file. You can scroll through the file or click on the **Search Next** or **Search Previous** buttons to search through the file. Click on the Cancel Search icon **x** at the top of the table to cancel the search.

You can also create and apply a filter to display specific log entries only. Click on the Filter icon  to bring up the Filter Selection window. Select an existing filter and click on the **Apply** button to apply the filter to the current display. You can also click on the **Add** icon to create a new filter. Enter a Filter Name and Filter Description, set the conditions for the filter and click on the **Add** button. Then select the filter to apply it to the current display. Click on the **Reset** button to cancel the filter and return to the unfiltered view of the file.

Click on the **Current** button (default) to display the contents of the current file. Click on the **Historical** button to display all ngnms.log file information, including archived files. Click on the **Browse All** button to view a list of all archived ngnms.log Files. Click on a file in the list to display the contents of the file. Click on the **Download** button at the top of the screen to save the file to your computer. Click on the Refresh icon  to refresh the display with the most recent log entries.

Viewing Log Files

By default, when you click on a log file (e.g., discovery.log, polling.log) the most current log entries are displayed. This "current" view is determined by the settings you configure on the Settings Screen. Once the number of entries reaches the configured number, the log file is archived with the current date and stored in OmniVista. When viewing a log file, the Audit application enables you to view the contents of the current selected log file, view the contents of all log files (current and archived) for the selected log

file, or a view list of archived files, which you can then open and view.

Click on a log to view contents.


Click **Download** to download the file. Click on the **Refresh** icon to display the most recent log entries.

Search through log entries.

Contents of selected log file.

To view a log file, click on a log file type the Home Page (e.g., Network, Configuration), or click on a log file type on the left side of the screen to display the list of log files for that type. Click on a log file (e.g., discovery.log, polling.log) to display the contents of the file. By default, the contents of the most current log file are displayed. This "current" view will display all of the contents that have not been archived. To view the contents of all log files for the selected log file (current and archived) click on the **Historical** button at the top of the screen.

To view a list of all log files (current and archived), click on the **Browse All** button at the top of the screen. A list of all log files for the selected log is displayed. The current log file is displayed at the top (e.g., polling.log) followed by all of the archived log files, identified by the date and time the file was archived (e.g., polling_04-29-2016_061727PM.rou). Click on a log file to display the contents of the file.

You can also download the contents of any log file you are viewing by clicking on the **Download** button at the top of the screen. The contents of the file can then be saved to your computer. Click on the Refresh icon  to refresh the display with the most recent log entries.

Searching Through Log Files

You can scroll through the contents of a log file using the scroll bar and arrows on the right side of the screen. You can also search a log file by keyword. Enter the search criteria in the "Search in this page" field at the top of the log file and click on the **Search Next** button. The word or phrase is highlighted in yellow throughout the file. You can scroll through the file or click on the **Search Next** or **Search Previous** buttons to search through the file. Click on the Cancel Search icon **x** at the top of the table to cancel the search.

Filtering Log File Entries

You can create and apply a filter to display specific log entries only. Click on the Filter icon to bring up the Filter Selection window. Select an existing filter and click on the **Apply** button to apply the filter to the current display. You can also click on the **Add** icon to create a new filter. Enter a Filter Name and Filter Description, set the conditions for the filter and click on the **Add** button. Then select the filter to apply it to the current display. Click on the **Reset** button to cancel the filter and return to the unfiltered view of

the file.

Downloading Log Files

You can download the contents of any log file you view by clicking on the **Download** button at the top of the screen. The contents of the log file will be downloaded as a text file that can be opened with any text editor.

All Current Logs

The Audit All Current Logs Screen displays a list of all current logs. Scroll or search through the list and click on a file to display the contents of the file. Select a file and click on the **Download** button to download the file to your computer. Click on the **Download All** button to download all of the log files to your computer. The files will be saved in a Zip file.

Settings

The Audit Settings Screen is used to specify Audit Log File preferences. Configure a field as described below and click the **Apply** button. The changes take effect immediately. If you change a field configuration, you can click on the **Revert** button to revert a field to the previous setting. Click on the **Default** button to set all of the fields to the default settings.

- **Maximum Audit Entries** - The maximum number of entries that can exist in any one log file (Range = 50 - 10,000, Default = 2,000). Any entry in excess of the value in this field will cause the current log file to be archived. For example, if the field is set to 1,000 entries, and a log file contains 1,000 entries, when the next entry is received the following will occur:
 - The current log file will be archived with 1,000 entries.
 - A new version of the file will be created that contains the latest entry.
- **Maximum Audit File Copies** - The maximum number of audit files that can be created. When the audit file reaches the configured maximum number of audit entries, the file is saved and a new file is started (Range = 0 - 100, Default = 5).
- **Max Log File Size** - The maximum log file size, in KB (Range = 1 - 30,000, Default = 10,240). Any entry in the file that increases the file size beyond the value in this field will cause current files to be archived. For example, if the field is set to 5120 KB, and a file is at a file size of 5120 KB, when the next entry is received the following will occur:
 - The current file will be archived with 5120 KB of information. The archive file will be located at *installation directory/data/logs*.
 - A new version of the file will be created that contains the latest entry.

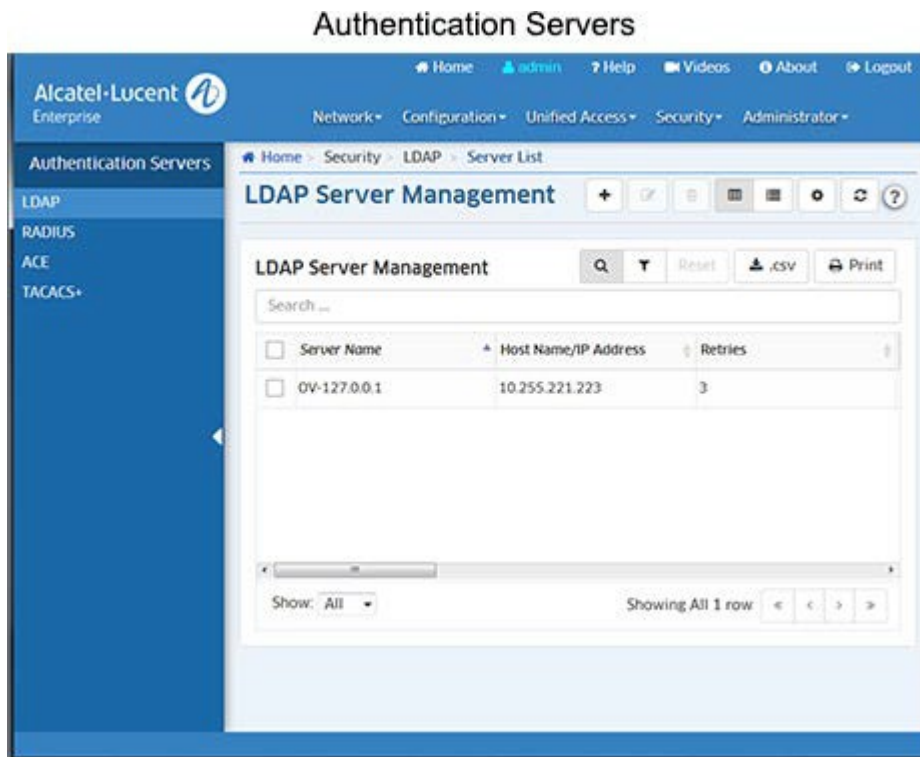
7.0 Authentication Servers

The Authentication Servers application enables you to create, modify, and delete authentication servers in OmniVista. An authentication server could be an LDAP, RADIUS, ACE, or TACACS+ Server. Any authentication server that you want to use, other than the default OmniVista LDAP Server, must be added to OmniVista. Adding a server to OmniVista basically informs OmniVista that the server exists. OmniVista does not search the network to locate available authentication servers, so any server that you add to OmniVista should actually exist (or should exist in the near future). When you add a server, you can also specify other information such as:

- Operating parameters for switches that will use the server for authentication, such as the number of retries the switch will attempt while communicating with the server.
- The user name and password used to login to the server (if applicable).
- The location of the server to be used as a "backup" server if the added server becomes unavailable.

Note: OmniVista cannot manage authentication server content.

When you open the Authentication Servers application, the LDAP Server Management Screen is displayed, and links to the LDAP, RADIUS, ACE, and TACACS+ screens are displayed on the left.



LDAP Servers

Lightweight Directory Access Protocol (LDAP) is a standard directory server protocol. The LDAP Server Management Screen lists all LDAP Authentication Servers known to OmniVista. It also enables you to add, modify, and delete LDAP Servers from the list of LDAP Servers known to OmniVista. By default, the OmniVista LDAP Server is automatically installed with OmniVista. However, any LDAP V3 server can be added to the list of known LDAP Servers. Click here for more information on configuring an LDAP Server.

RADIUS Servers

Remote Authentication Dial-in User Service (RADIUS) is a standard authentication and accounting protocol defined in RFC 2865 and RFC 2866. A built-in RADIUS Client is available in Alcatel-Lucent Enterprise AOS Switches. A RADIUS Server that supports Vendor Specific Attributes (VSAs) is required. VSAs carry specific authentication, authorization, and configuration details about RADIUS requests to and replies from the server. The RADIUS Server Management Screen lists all RADIUS Authentication Servers known to OmniVista. It also enables you to add, modify, and delete Servers from the list of RADIUS Servers known to OmniVista. [Click here for more information on configuring an RADIUS Server.](#)

ACE Servers

You can use a single external ACE Server for authentication of all switch access types. You are limited to a single ACE Server, because file **sdconf.rec** must be FTPed from the ACE Server to the switch's **/network** directory, to inform the switch of the ACE Server's IP address and other configuration information. This requirement means that the switch can communicate with only a single ACE Server at any one time. The ACE Server Management Screen enables you to add a single ACE Server to OmniVista. It also enables you to delete an ACE Server. An ACE Server cannot be configured or modified from OmniVista because all configuration information is contained in **sdconf.rec** file. [Click here for more information on configuring an ACE Server.](#)

TACACS+ Servers

Terminal Access Controller Access Control System (TACACS+) is a standard authentication and accounting protocol defined in RFC 1321 that employs TCP for reliable transport. A built-in TACACS+ Client is available in the switch. A TACACS+ Server allows access control for routers, network access servers, and other networked devices through one or more centralized servers. The protocol also allows separate authentication, authorization, and accounting services. By allowing arbitrary length and content authentication exchanges, it allows clients to use any authentication mechanism. The TACACS+ Server Management Screen lists all TACACS+ Authentication Servers known to OmniVista. It also enables you to add and, modify, and delete Servers from the list of TACACS+ Servers known to OmniVista. [Click here for more information on configuring a TACACS+ Server.](#)

LDAP Server Management

The Authentication Servers LDAP Server Management Screen displays all LDAP Authentication Servers known to OmniVista. It also enables you to add, modify, and delete LDAP Servers from the list of LDAP Servers known to OmniVista. Lightweight Directory Access Protocol (LDAP) is a standard directory server protocol. The LDAP Client in the switch is based on several RFCs: 1798, 2247, 2251, 2252, 2253, 2254, 2255, and 2256. The protocol was developed as a way to use directory services over TCP/IP and to simplify the Directory Access Protocol (DAP) defined as part of the Open Systems Interconnection (OSI) effort. Originally, LDAP was a front-end for X.500 DAP.


The LDAP protocol synchronizes and governs the communications between the LDAP Client and the LDAP Server. The protocol also dictates how database information, which is normally stored in hierarchical form, is searched from the root directory down to distinct entries. In addition, LDAP has its own format that permits LDAP-enabled Web browsers to perform directory searches over TCP/IP.

The OmniVista LDAP Server is automatically installed in OmniVista. You cannot modify or delete it. If you choose to use the OmniVista LDAP Server, you can easily add, modify, or delete users and user privileges in the server's LDAP database. However, if you want to use a different LDAP V3 server, you must add it to OmniVista. OmniVista only manages the built-in LDAP Server, other authentication servers must be managed outside of OmniVista. You can assign switches to such servers, but the Authentication Servers

application does not allow you to add, modify, or delete users and user privileges in the LDAP database of such servers. This is because an LDAP Server's database must be configured for the specific schema used to manage users and there is no public API for configuring LDAP schemas. Click here for more information on configuring an LDAP Server.

Note: LDAP Server Management supports both AOS and wireless devices; however certain attributes may not be supported on wireless devices. See the configuration fields below for more information.

Adding an LDAP Server


As mentioned earlier, the OmniVista LDAP Server is automatically installed in OmniVista and known to OmniVista. However, if you have configured a new LDAP Server, you must add it to the list of LDAP Servers known to OmniVista. To add a new LDAP Server, click on the Create icon  and complete the fields as described below. When you are finished, click the **Create** button.

- **Server Name** - A unique name for the LDAP Authentication Server. This name will be used by OmniVista and the switch to identify the server.
- **Host Name/IP Address** - The name of the computer where the server is located OR the IP address of the computer where the server is located.
- **Backup Host Name/IP Address** - Each LDAP Server may optionally have a backup server. If you wish to define a backup server that will be used if this server is unavailable, enter the name of the computer where the backup server is located OR enter the IP address of the computer where the backup server is located.
- **Retries** - The number of retries that you want the switch to attempt when trying to contact the LDAP Server (Range = 1 - 3, Default = 3). (Not supported on wireless devices and ignored when applied to those devices.)
- **Timeout** - The number of seconds that you want the switch to wait before a request to the LDAP authentication server is timed out (Range = 1 - 30, Default = 2).
- **Password** - Password used to login to the LDAP Server.
- **Confirm Password** - Re-enter the LDAP Server password.
- **SSL** - Set this field to **True** or **False** to inform the switch whether SSL (Secure Socket Layer) is enabled or disabled on the LDAP authentication server. SSL can be set up on the server for additional security. (This usually involves adding digital certificates to the server.) When SSL is enabled, the server's identity will be authenticated. Refer to "Managing Authentication Servers" in your *Network Configuration Guide* and to the instructions provided by the LDAP Server's vendor for further information on setting up SSL on the LDAP Server. (Not supported on wireless devices and ignored when applied to those devices.)
- **Port** - The port number used as the LDAP port address. This is the port at which the LDAP Server "listens". By default, the port number is 389. However, note that the switch automatically sets the port number to 636 when SSL is enabled. (Port number 636 is typically used on LDAP Servers for SSL.) The port number on the switch must match the port number configured on the server.
- **Admin Name** - The name used to login to the LDAP Server.
- **Search Base** - The search base in the LDAP Server where authentication information can be found (e.g., o=alcatel.com).
- **VRF Name** - The VRF Instance associated with the LDAP Server, if applicable. An LDAP Server can be configured on any VRF instance including the default VRF instance. However, all of the servers (for example, all the LDAP servers) must reside on the same VRF instance. Default value is "default". Note that VRF Name is not supported on wireless devices and will be ignored when applied to those devices.

Note: SSL communication with the LDAP Server is not supported on OS6860 Switches (AOS


8.1.1 R01).

Modifying an LDAP Server

Select an LDAP Server in the list and click on the Edit  icon. Edit any necessary fields as described above, then click on the **Save** button. It is important to note that you cannot modify values indiscriminately. The values must match those of the actual LDAP Server. For example, if you want to change the LDAP port address, you must first use the tools provided by your LDAP Server's vendor to change the port on the LDAP Server itself. You can then inform OmniVista that the port number has changed by modifying the **Port** field. Also note that you cannot edit an LDAP Server's name. You must delete it and create a new one.

Note: You cannot delete an LDAP Server that is currently being used by OmniVista.

Deleting an LDAP Server

Select an LDAP Server in the list and click on the Delete icon . Note that deleting an LDAP server will not cause switches that currently use that server to cease using it. Switches using the deleted LDAP Server will continue to use it until the switches are reassigned.

Configuring an LDAP Server

As mentioned earlier, the OmniVista LDAP Server is automatically installed along with the authentication server. However, if you want to use a different LDAP V3 server, you must add it to OmniVista. You can assign switches to such servers, but the Authenticated Servers application does not allow you to add, modify, or delete users and user privileges in the LDAP database of such servers. This is because an LDAP Server's database must be configured for the specific schema used to manage users and there is no public API for configuring LDAP schemas.

Before you add an LDAP Server to OmniVista's list of available authentication servers, you must first install the LDAP Server based on the instructions provided by the LDAP Server's vendor. You must then modify the LDAP Server's schema to add the LDAP objects required to manage Alcatel-Lucent Enterprise Switches, and configure user accounts on the server.

Required LDAP Schema Objects

The following objects must be added to an LDAP Server's schema so that it can manage Alcatel-Lucent Enterprise Switches. To modify the schema, follow the vendor's instructions. Each LDAP vendor provides a different way of modifying the schema.

- attribute accountfailtime oid-ataccountfailtime cis
- attribute accountstarttime oid-ataccountstarttime cis
- attribute accountstoptime oid-ataccountstoptime cis
- attribute numberofswitchgroups oid-atnumberofswitchgroups int single
- attribute switchgroups oid-atswitchgroups int
- attribute switchserialnumber oid-atswitchserialnumber cis
- attribute switchslotport oid-atswitchslotport cis
- attribute clientipaddress oid-atclientipaddress cis
- attribute clientmacaddress oid-atclientmacaddress cis
- attribute userPermissions oid-atuserPermissions int single
- attribute pm-access-priv-read-1 oid-atpm-access-priv-read-1 cis single
- attribute pm-access-priv-read-2 oid-atpm-access-priv-read-2 cis single
- attribute pm-access-priv-write-1 oid-atpm-access-priv-write-1 cis single
- attribute pm-access-priv-write-2 oid-atpm-access-priv-write-2 cis single

- attribute pm-access-priv-global-1 oid-atpm-access-priv-global-1 cis single
- attribute pm-access-priv-global-2 oid-atpm-access-priv-global-2 cis single
- attribute bop-asa-func-priv-read-1 oid-atbop-asa-func-priv-read-1 int single
- attribute bop-asa-func-priv-read-2 oid-atbop-asa-func-priv-read-2 int single
- attribute bop-asa-func-priv-write-1 oid-atbop-asa-func-priv-write-1 int single
- attribute bop-asa-func-priv-write-2 oid-atbop-asa-func-priv-write-2 int single
- attribute allowedTime oid-atallowedTime cis single
- attribute bop-asa-geo-priv-profile-number oid-atbop-asa-geo-priv-profile-number int single
- attribute bop-md5key oid-atbop-md5key cis single
- attribute bop-shakekey oid-atbop-shakekey cis single
- attribute bop-asa-snmp-level-security oid-atbop-asa-snmp-level-security int single

Configuring User Accounts on the Server

When you use an LDAP Server other than the OmniVista LDAP Server, you must set up all user accounts on the server based on the instructions provided by the LDAP Server's vendor. You cannot set up user accounts from OmniVista for any authentication server other than the OmniVista LDAP Server, which is automatically installed with the Authentication Servers application.

LDAP Server Management Table

The LDAP Server Management Table displays information about all LDAP Servers known to OmniVista.

- **Server Name** - A unique name for the LDAP Authentication Server. This name will be used by OmniVista and the switch to identify the server.
- **Host Name/IP Address** - The name of the computer where the server is located OR the IP address of the computer where the server is located.
- **Retries** - The number of retries that you want the switch to attempt when trying to contact the LDAP Server (Range = 1 - 3, Default = 3). (Not supported on wireless devices and ignored when applied to those devices.)
- **Timeout** - The number of seconds that you want the switch to wait before a request to the LDAP authentication server is timed out (Range = 1 - 30, Default = 2).
- **Admin Name** - The name used to login to the LDAP Server.
- **Search Base** - The search base in the LDAP Server where authentication information can be found (e.g., o=alcatel.com).
- **Port** - The port number used as the LDAP port address. This is the port at which the LDAP Server "listens". By default, the port number is 389. However, note that the switch automatically sets the port number to 636 when SSL is enabled. (Port number 636 is typically used on LDAP Servers for SSL.) The port number on the switch must match the port number configured on the server.
- **SSL** - Set this field to **True** or **False** to inform the switch whether SSL (Secure Socket Layer) is enabled or disabled on the LDAP authentication server. SSL can be set up on the server for additional security. (This usually involves adding digital certificates to the server.) When SSL is enabled, the server's identity will be authenticated. Refer to "Managing Authentication Servers" in your *Network Configuration Guide* and to the instructions provided by the LDAP Server's vendor for further information on setting up SSL on the LDAP Server. (Not supported on wireless devices and ignored when applied to those devices.)
- **Backup Host Name/IP Address** - Each LDAP Server may optionally have a backup server. If you wish to define a backup server that will be used if this server is unavailable, enter the name of the computer where the backup server is located OR enter the IP address of the computer where the backup server is located.
- **VRF Name** - The VRF Instance associated with the LDAP Server, if applicable. An LDAP Server can be configured on any VRF instance including the default VRF instance. However, all of the servers (for

example, all the LDAP servers) must reside on the same VRF instance. Default value is "default". Note that VRF Name is not supported on wireless devices and will be ignored when applied to those devices.

RADIUS Server Management

The Authentication Servers RADIUS Server Management Screen displays all RADIUS Servers known to OmniVista. It also enables you to add, modify, and delete RADIUS Servers from the list of RADIUS Servers known to OmniVista. A built-in RADIUS Client is available in the switch. A RADIUS Server that supports Vendor Specific Attributes (VSAs) is required. VSAs carry specific authentication, authorization, and configuration details about RADIUS requests to and replies from the server. Refer to "Managing Authentication Servers" in your *Network Configuration Guide* for specific information on the VSAs required. Before you add a RADIUS Server to OmniVista's list of RADIUS Servers known to OmniVista, you must first install and configure the RADIUS Server.

Note: You cannot add, modify, or delete users and user privileges from RADIUS Servers in OmniVista.


Note: RADIUS Server Management supports both AOS and wireless devices; however certain attributes may not be supported on wireless devices. See the configuration fields below for more information.

Adding a RADIUS Server


After configuring a RADIUS Server, you must add it to the list of RADIUS Servers known to OmniVista. To add a new RADIUS Server, click on the Create icon **+** and complete the fields as described below. When you are finished, click the **Create** button.

- **Server Name** - Unique name for the RADIUS Server. This name will be used by OmniVista and the switch to identify the Server.
- **Host Name/IP Address** - The name of the computer where the server is located OR the IP address of the computer where the Server is located.
- **Backup Host Name/IP Address** - Each RADIUS Server may optionally have a backup server. If you wish to define a backup server that will be used if this server is unavailable, enter the name of the computer where the backup server is located OR enter the IP address of the computer where the backup Server is located. (Not supported on wireless devices and ignored when applied to those devices.)
- **Retries** - The number of retries that you want the switch to attempt when trying to contact the RADIUS Server (Range = 1 - 3, Default = 3).
- **Timeout** - The number of seconds that you want the switch to wait before a request to the RADIUS Server is timed out (Range = 1 - 30, Default = 2).
- **Shared Secret** - The password to the Server. (The "shared secret" is essentially the server password.) Note that the password you enter must be configured identically on the Server.
- **Confirm Secret** - Re-enter the Shared Secret.
- **Authentication Port** - The port you to access the Server (Range = 1 - 65535, Default = 1812).
- **Accounting Port** - The port for accounting information (Range = 1 - 65535, Default = 1813).
- **VRF Name** - The VRF Instance associated with the RADIUS Server, if applicable. A RADIUS Server can be configured on any VRF instance including the default VRF instance. However, all of the servers (for example, all the RADIUS servers) must reside on the same VRF instance. Default value is "default". (Not supported on wireless devices and ignored when applied to those devices.)

Modifying a RADIUS Server

Select a RADIUS Server in the list and click on the Edit  icon. Edit any necessary fields as described above, then click on the **Save** button. It is important to note that you cannot modify values indiscriminately. The values must match those of the actual RADIUS Server. For example, if you want to change the RADIUS Authentication port, you must first use the tools provided by your RADIUS Server's vendor to change the port on the RADIUS Server itself. You can then inform OmniVista that the port number has changed by modifying the **Authentication Port** field.

Deleting a RADIUS Server

Select a RADIUS Server in the list and click on the Delete icon . Note that deleting an authentication server from the list of RADIUS Servers known to OmniVista will not cause switches that currently use that RADIUS Server to cease using it. Switches using the deleted RADIUS Server will continue to use it until the switches are reassigned.

Configuring a RADIUS Server

Before you add a RADIUS Server to OmniVista's list of RADIUS Servers known to OmniVista, you must first install and configure the RADIUS Server, then configure user accounts on the Server.

Configuring the Server

Before you add a RADIUS Server to OmniVista's list of available authentication servers, you must first install the RADIUS Server based on the instructions provided by the RADIUS Server's vendor. Then, you must configure the RADIUS Server with the vendor specific attributes. These attributes carry specific authentication, authorization, and configuration details about RADIUS requests to and replies from the server. Refer to "Managing Authentication Servers" in your Network Configuration Guide for specific information on the VSAs required.

Configuring User Accounts on the Server

When you use a RADIUS Server for User Authentication, you must set up all user accounts on the server based on the instructions provided by the RADIUS Server's vendor. However, the authorization which includes the access level associated with each user will be controlled by the OmniVista server, based on the group a user is associated with. You cannot set up user accounts from OmniVista for any authentication server other than the OmniVista LDAP server, which is automatically installed with the Authentication Servers application.

Once you have installed, configured, and set up the user accounts on the RADIUS Server, you are ready to add the server to OmniVista.

RADIUS Server Management Table

The RADIUS Server Management Table displays information about all RADIUS Servers known to OmniVista.


- **Server Name** - Unique name for the RADIUS Server. This name will be used by OmniVista and the switch to identify the Server.
- **Host Name/IP Address** - The name of the computer where the server is located OR the IP address of the computer where the Server is located.
- **Retries** - The number of retries that you want the switch to attempt when trying to contact the RADIUS Server (Range = 1 - 3, Default = 3).

- **Timeout** - The number of seconds that you want the switch to wait before a request to the RADIUS Server is timed out (Range = 1 - 30, Default = 2).
- **Authentication Port** - The port you to access the Server (Range = 1 - 65535, Default = 1812).
- **Accounting Port** - The port for accounting information (Range = 1 - 65535, Default = 1813).
- **VRF Name** - The VRF Instance associated with the RADIUS Server, if applicable. A RADIUS Server can be configured on any VRF instance including the default VRF instance. However, all of the servers (for example, all the RADIUS servers) must reside on the same VRF instance. Default value is "default". (Not supported on wireless devices and ignored when applied to those devices.)
- **Backup Host Name/IP Address** - Each RADIUS Server may optionally have a backup server. If you wish to define a backup server that will be used if this server is unavailable, enter the name of the computer where the backup server is located OR enter the IP address of the computer where the backup Server is located. (Not supported on wireless devices and ignored when applied to those devices.)


ACE Server Management

The Authentication Servers ACE Server Management Screen can be used to add or delete an ACE Server. You can use a single external ACE Server for authentication of all switch access types. You are limited to a single ACE Server, because file **sdconf.rec** must be FTPed from the ACE Server to the switch's **/network** directory, to inform the switch of the ACE Server's IP address and other configuration information. This requirement means that the switch can communicate with only a single ACE Server at any one time. The ACE Server Management Screen enables you to add a single ACE Server to OmniVista. It also enables you to delete an ACE Server. An ACE Server cannot be configured or modified from OmniVista because all configuration information is contained in **sdconf.rec** file. Note that an ACE Server cannot be used for Layer 2 authentication or for policy.

Adding an ACE Server

Once you have installed and configured the ACE Server, you must add it to the list of ACE Servers known to OmniVista. To add a new ACE Server, click on the Create icon . When you assign the ACE Server to switches, the authentication server will automatically configure the switches to operate with the server.

Deleting an ACE Server

To delete an ACE Server from OmniVista, select the Server and click on the Delete icon . Deleting the ACE Servers known to OmniVista will not cause switches that currently use that ACE Server to cease using it. Switches using the deleted ACE Server will continue to use it until the switches are reassigned.

Configuring an ACE Server

Before you add the ACE Server to OmniVista, you must first install the ACE Server, based on the instructions provided by your ACE Server's vendor. You must also set up user accounts on the ACE Server. There are no server-specific parameters that must be configured for the switch to communicate with an attached ACE Server; however, you must FTP the **sdconf.rec** file from the server to the switch's **/network** directory. This file is required so that the switch will know the IP address of the ACE Server and other configuration information. For information about loading files into the switch, see the *OmniSwitch Switch Management Guide*.

Note: An ACE Server stores and authenticates switch user accounts (i.e., user IDs and passwords), but does NOT store or send user privilege information to the switch. User privileges for logins are determined by the switch itself. When a user attempts to log into the switch, the user ID and password are sent to the ACE Server. The server determines whether the login is valid or not. If the login is valid, the user privileges must be determined. The switch checks its

user database for the user's privileges. If the user is not in the database, the switch uses the default privilege, which is determined by the default user account. For information about the default user account, see the "Switch Security" chapter of the *OmniSwitch Switch Management Guide*.

The ACE client in the switch is version 4.1; it does not support the replicating and locking feature of ACE 5.0, but it may be used with an ACE 5.0 server if a legacy configuration file is loaded on the server. The legacy configuration must specify authentication to two specific servers (master and slave). See the RSA Security ACE Server documentation for more information.

TACACS+ Server Management

The Authentication Servers TACACS+ Server Management Screen displays all TACACS+ Authentication Servers known to OmniVista. It also enables you to add, modify, and delete TACACS+ Servers from the list of TACACS+ Servers known to OmniVista. Terminal Access Controller Access Control System (TACACS+) is a standard authentication and accounting protocol defined in RFC 1321 that employs TCP for reliable transport. A built-in TACACS+ Client is available in the switch. A TACACS+ Server allows access control for routers, network access servers, and other networked devices through one or more centralized servers. The protocol also allows separate authentication, authorization, and accounting services. By allowing arbitrary length and content authentication exchanges, it allows clients to use any authentication mechanism.

The TACACS+ Client offers the ability to configure multiple TACACS+ Servers. When the primary server fails, the client tries the subsequent servers. Multiple server configurations are applicable only for backup and not for server chaining.

Note: TACACS+ Server Management supports both AOS and wireless devices; however certain attributes may not be supported on wireless devices. See the configuration fields below for more information.


Adding a TACACS+ Server

Once you have installed, configured, and set up the user accounts on the TACACS+ Server, you are ready to add the server to OmniVista. To add a new TACACS+ Server, click on the Create icon **+** and complete the fields as described below. When you are finished, click the **Create** button.


- **Server Name** - A unique name for the TACACS+ Authentication Server. This name will be used by OmniVista and the switch to identify the server.
- **Host Name/IP Address** - The name of the computer where the server is located OR the IP address of the computer where the server is located.
- **Backup Host Name/IP Address** - Each TACACS+ Server may optionally have a backup server. If you wish to define a backup server that will be used if this server is unavailable, enter the name of the computer where the backup server is located OR enter the IP address of the computer where the backup server is located.
- **Timeout** - The number of seconds that you want the switch to wait before a request to the TACACS+ authentication server is timed out. Default value is 2.
- **Shared Secret** - The password to the Server. (The "shared secret" is essentially the server password.) Note that the password you enter must be configured identically on the Server.
- **Confirm Shared Secret** - Re-enter the Shared Secret.
- **Port** - The port number used to access the TACACS+ Server (Default = 49).
- **VRF Name** - The VRF Instance associated with the TACACS+ Server, if applicable. A TACACS+ Server can be configured on any VRF instance including the default VRF instance. However, all of the servers (for example, all the TACACS+ servers) must reside on the same VRF instance. Default

value is "default". (Not supported on wireless devices and ignored when applied to those devices.)

Modifying a TACACS+ Server

Select a TACACS+ Server in the list and click on the Edit  icon. Edit any necessary fields as described above, then click on the **Save** button. It is important to note that you cannot modify values indiscriminately. The values must match those of the actual TACACS+ Server. For example, if you want to change the TACACS+ authentication port, you must first use the tools provided by your TACACS+ Server's vendor to change the port on the TACACS+ Server itself. You can then inform OmniVista that the port number has changed by modifying the **Port** field.

Deleting a TACACS+ Server

Select a TACACS+ Server in the list and click on the Delete icon . Note that deleting a TACACS+ Server will not cause switches that currently use that TACACS+ Server to cease using it. Switches using the deleted TACACS+ Server will continue to use it until the switches are reassigned.

TACACS+ Server Management Table

The TACACS+ Server Management Table displays information about all TACACS+ Authentication Servers known to OmniVista.

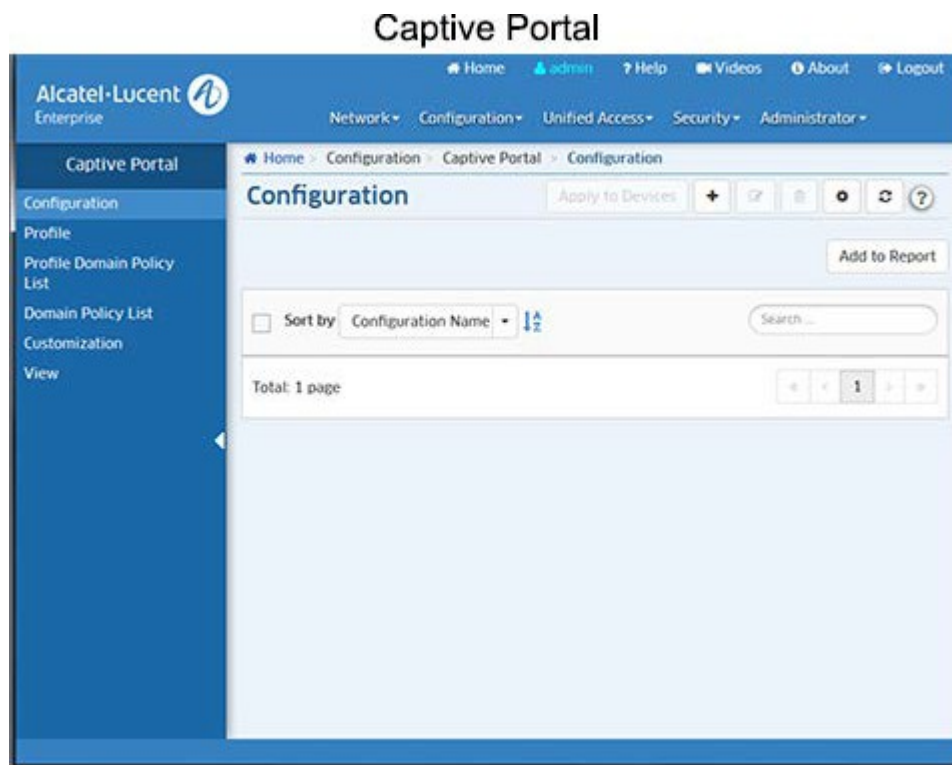
- **Server Name** - A unique name for the TACACS+ Authentication Server. This name will be used by OmniVista and the switch to identify the server.
- **Host Name/IP Address** - The name of the computer where the server is located OR the IP address of the computer where the server is located.
- **Timeout** - The number of seconds that you want the switch to wait before a request to the TACACS+ authentication server is timed out. Default value is 2.
- **Port** - The port number used to access the TACACS+ Server (Default = 49).
- **VRF Name** - The VRF Instance associated with the TACACS+ Server, if applicable. A TACACS+ Server can be configured on any VRF instance including the default VRF instance. However, all of the servers (for example, all the TACACS+ servers) must reside on the same VRF instance. Default value is "default". (Not supported on wireless devices and ignored when applied to those devices.)
- **Backup Host Name/IP Address** - Each TACACS+ Server may optionally have a backup server. If you wish to define a backup server that will be used if this server is unavailable, enter the name of the computer where the backup server is located OR enter the IP address of the computer where the backup server is located.

8.0 Captive Portal

Captive Portal authentication is mechanism by which user credentials are obtained through Web pages and authenticated through a RADIUS server. If the authentication is successful, the RADIUS server may return a role (policy list) that is applied to traffic from the user device. The OmniSwitch implementation supports an internal Captive Portal mechanism. An internal Web server on the local switch presents Captive Portal Web pages to obtain user credentials.

Internal Captive Portal authentication is a configurable option for a UNP Access Role Profile that is applied after a user is assigned to the profile (after the initial 802.1X or MAC authentication or classification process). This type of authentication does not change the Access Role Profile assignment for the user device. Instead, Captive Portal provides a secondary level of authentication that is used to apply a new role (QoS policy list) to the user. An external, guest Captive Portal authentication mechanism is provided through the Bring Your Own Device (BYOD) feature, which integrates Access Guardian with ClearPass Policy Manager (CPPM).

For more information, see the "Using Captive Portal Authentication" section of the Access Guardian chapter in the *OmniSwitch Network Configuration Guide*.




Configuration

The Captive Portal Configuration Screen displays all configured Captive Portal global configurations and is used to create, edit, and delete Captive Portal global Configurations.


Creating a Captive Portal Configuration

Click on the Create icon **+**. Enter a **Configuration Name**, configure the fields as described below, then click on the **Create** button. When you are finished, select the checkbox next to the profile and click on the **Apply to Devices** button to assign the configuration to switches on the network.

Configuration

- **Configuration Name** - The unique name for the Configuration. This name will be used by OmniVista and not to be assigned to the switch.
- **Redirect IP Address** - The IP address of the Redirect Server.
- **Redirect Host/DNS Name** - The Host Name/DNS Name of the Redirect Server.
- **Proxy Port** - The TCP port on the proxy server.
- **Success Redirect URL** - The Captive Portal's Redirect URL upon successful authentication. **Retries** - The number of times a device can try to login before Captive Portal determines that authentication for that device has failed (Range = 1 - 99, Default = 3).
- **Policy List** - The Unified Policy List to apply to the authenticated user device. (You can also click on the Add icon  to go to the Unified Policy List Screen and create a Policy List.)

Editing a Captive Portal Configuration


Select a Configuration in the list and click on the Edit icon  to bring up the Edit Captive Portal Configuration Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server. Note that you cannot edit the Configuration Name.

Assigning a Captive Portal Configuration

When you click the **Apply to Devices** button, the Apply to Devices Screen appears. Select the switch(es) to which you want to assign the configuration and click on the **Apply** button. Click **OK** to return to the Configuration Screen.

Note: To "unassign" a configuration from a switch(es), select the configuration and click on the **Apply to Devices** button. The Apply to Devices Screen appears. The switches to which the configuration is assigned will appear on the right. Move the switches from which you want to "unassign" the configuration to the left side of the screen and click the **Apply** button.


Deleting a Captive Portal Configuration

Select a Configuration in the list and click on the Delete icon , then click **OK**, at the confirmation prompt. If the configuration has been assigned to a switch, a second prompt will appear. Click **OK** on the second prompt to delete the configuration.



Profile

The Captive Portal Profile Screen displays all configured Captive Portal Profiles and is used to create, edit, and delete Captive Portal Profiles. A Captive Portal profile is a configuration entity that provides flexible assignment of Captive Portal configuration parameters to devices classified into specific UNP Access Role Profiles. However, this type of profile is only valid when assigned to Access Role Profiles on which Captive Portal authentication is enabled. When a Captive Portal profile is applied to a UNP Access Role Profile, the parameter values defined in the Captive Portal profile override the global Captive Portal parameter values configured for the switch. If there is no Captive Portal profile associated with an Access Role Profile, the global Captive Portal configuration is applied.


Creating a Captive Portal Profile

Click on the Create icon . Enter a **Profile Name**, configure the fields as described below, then click on the **Create** button. When you are finished, select the checkbox next to the profile and click on the **Apply to Devices** button to assign the profile to switches on the network.

Captive Portal Profile Configuration

- **Profile Name** - The unique name for the Captive Portal profile.
- **Policy List** - The QoS policy list to apply when Captive Portal authentication is successful but the RADIUS server did not return a policy list. (You can also click on the Add icon  to go to the Unified Policy List Screen and create a Policy List.)
- **AAA Server Profile** - The AAA profile used to define specific device authentication configuration options, such as which servers to use for Captive Portal authentication and parameter values for session timers and RADIUS attributes. If there is no AAA profile assigned, the global AAA configuration for the switch is used. (You can also click on the Add icon  to go to the Unified Policy List Screen and create a AAA Server Profile.)
- **Success Redirect URL Name** - The Captive Portal's Redirect URL upon successful authentication.
- **Retries** - Number of Captive Portal retries before failure is declared.

Editing a Captive Portal Profile


Select a Captive Portal Profile in the list and click on the Edit icon  to bring up the Edit Captive Portal Profile Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server. Note that you cannot edit the Profile Name.

Assigning a Captive Portal Profile

When you click the **Apply to Devices** button, the Apply to Devices Screen appears. Select the switch(es) to which you want to assign the Captive Portal Profile and click on the **Apply** button. Click **OK** to return to the Configuration Screen.

Note: To "unassign" a profile from a switch(es), select the profile and click on the **Apply to Devices** button. The Apply to Devices Screen appears. The switches to which the profile is assigned will appear on the right. Move the switches from which you want to "unassign" the profile to the left side of the screen and click the **Apply** button.

Deleting a Captive Portal Profile

Select a Captive Portal Profile in the list and click on the Delete icon , then click **OK**, at the confirmation prompt. If the profile has been assigned to a switch, a second prompt will appear. Click **OK** on the second prompt to delete the profile.

Profile Domain Policy List

The Captive Portal Profile Domain Policy List Screen displays all configured Captive Portal Domain Profiles and is used to create, edit, and delete Captive Portal Domain Profiles. A Captive Portal Domain Profile is used to assign a Captive Portal Profile and QoS Policy List to users logging in from a specific domain (e.g. NA02/tut).


Creating a Profile Domain Policy

Click on the Create icon **+** . Select a Captive Portal **Profile Name** from the drop-down list. (You can also click on the Add icon **+** to go to the Captive Portal Profile Screen to create a profile.) Configure the fields as described below, then click on the **Create** button. When you are finished, select the checkbox next to the profile and click on the **Apply to Devices** button to assign the configuration to switches on the network.

Profile Domain Pass Policy Configuration

- **Profile Name** - Select a Captive Portal Profile from the drop-down menu. (You can also click on the Add icon **+** to go to the Captive Portal Profile Screen to create a profile.)
- **Domain** - A unique name for the Captive Portal Domain Policy.
- **Policy List** - Select a Unified Policy List from the drop-down menu. This Policy List will replace the one in the Access Role Profile upon successful Captive Portal authentication only if the domain is known. (You can also click on the Add icon **+** to go to the Unified Policy List Screen and create a Policy List.)
- **Realm** - Select Suffix/Prefix. For example: Suffix: NA02/tut; Prefix: tu@alu.com.

Editing a Profile Domain Policy


Select a policy and click on the Edit icon  to bring up the Edit Profile Domain Policy List Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server. Note that you cannot edit the Profile Name.

Assigning a Profile Domain Policy

When you click the **Apply to Devices** button, the Apply to Devices Screen appears. Select the switch(es) to which you want to assign the policy and click on the **Apply** button. Click **OK** to return to the Configuration Screen.

Note: To "unassign" a profile from a switch(es), select the policy and click on the **Apply to Devices** button. The Apply to Devices Screen appears. The switches to which the profile is assigned will appear on the right. Move the switches from which you want to "unassign" the profile to the left side of the screen and click the **Apply** button.

Deleting a Profile Domain Policy

Select a policy and click on the Delete icon , then click **OK**, at the confirmation prompt. If the profile has been assigned to a switch, a second prompt will appear. Click **OK** on the second prompt to delete the profile.


Domain Policy List

The Captive Portal Domain Policy List Screen displays all configured Captive Portal Domain Policy Lists to create, edit, and delete Captive Portal Policy Lists. This screen enables you to define Policy Lists for different realms in which the endpoints are successfully authenticated. This is similar to creating a Profile Domain Policy List without the profile coming into play.


Creating a Captive Portal Domain Policy List

Click on the Create icon **+** . Enter a **Profile Name**, configure the fields as described below, then click on the **Create** button. When you are finished, select the checkbox next to the profile and click on the **Apply to Devices** button to assign the profile to switches on the network.

Captive Portal Domain Policy List Configuration

- **Domain** - The Captive Portal Profile Authenticated Domain Name.
- **Policy List** - The Unified Policy List to apply when Captive Portal authentication is successful but the RADIUS server did not return a policy list. (You can also click on the Add icon  to go to the Unified Policy List Screen and create a Policy List.)
- **Realm** - The Captive Portal Profile Authenticated Pass Realm.

Editing a Captive Portal Domain Policy List


Select a Policy List and click on the Edit icon  to bring up the Edit Domain Policy List Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server. Note that you cannot edit the Domain.

Assigning a Captive Portal Domain Policy List

When you click the **Apply to Devices** button, the Apply to Devices Screen appears. Select the switch(es) to which you want to assign the Policy List and click on the **Apply** button. Click **OK** to return to the Configuration Screen.

Note: To "unassign" a Policy List from a switch(es), select the profile and click on the **Apply to Devices** button. The Apply to Devices Screen appears. The switches to which the profile is assigned will appear on the right. Move the switches from which you want to "unassign" the profile to the left side of the screen and click the **Apply** button.

Deleting a Captive Portal Domain Policy List


Select a Policy List and click on the Delete icon , then click **OK**, at the confirmation prompt. If the profile has been assigned to a switch, a second prompt will appear. Click **OK** on the second prompt to delete the Policy List.



Captive Portal Customization

The Captive Portal Customization Screen displays all Customized Captive Portal web page files and is used to create, edit, and delete custom Captive Portal web page files. These files (e.g., html files, jpeg files) are used to create the web pages that are presented to the user during Captive Portal Login. OmniSwitches contain a default set of Captive Portal web page files. However, OmniVista enables you to import and then customize these files. You can then upload these custom files to switches. When you import the files from a switch using OmniVista, all of the necessary Captive Portal web page files are zipped together in an "archive" file. You can then customize the Captive Portal web pages by editing individual files within this archive and then uploading the edited archive to a switch(es). When a customized archive is uploaded to a switch, Captive Portal presents these web pages to the user, rather than the default pages stored on the switch.

Creating a Captive Portal Customization

Click on the Create icon  to bring up the Customization Workflow Wizard. Creating a Custom Captive Portal Customization consists of the following steps. After completing the steps as described below, click the **Create** button.

1. Import the Archive File
2. Download the Archive File for Editing
3. Upload the Edited Archive File
4. Apply the Archive File to Switches

Import the Archive

Enter a **File Set Name** for your customized set of files and an optional **Description**, then import the files. As mentioned earlier, OmniVista zips all of the Captive Portal web page files into a single archive file). You can actually import the file from a switch or import an existing archive from a local machine.

- **Import from a Switch** - Click on the **Select Switch** button and select a switch from which to download the files. Select **Release Folder** to import the default set of Captive Portal web page files or **Custom Folder** to select a previous set of customized files, then click on the **Import** button. Click on the **Next** button to download the file.
- **Import from Local Machine** - If you already have an edited archive that you want to upload from a local machine, click on the **Browse** button to locate the file. Click **Upload** to upload the file to OmniVista, then click **Next** to apply the file to switches. Note that if you import the file from a local machine, Steps 2 and 3 will be skipped and you only need to apply the edited file to switches.

Download the Archive File for Editing

If you import the Archive File from a switch, click on **Download** button then click on the download prompt to download the file. The file will be downloaded to your default download directory. Edit the necessary file(s) in the archive using the editor of your choice. When you are done, click on the **Next** button to upload the file.


Upload the Edited File

Click on the **Browse** button to locate the edited Archive File, then click on the **Upload** button to upload the file to OmniVista, then click the **Upload** button.

Apply the File to Switches

Select the switch(es) to which you want to assign the file and click on the **Finish** button. Note that you can also apply existing Archives to switches by selecting the Archive in the list and clicking on the **Push to Switches** button. The new Custom Archive that you apply will replace any existing Custom Archive on the switch(es).

Editing a Captive Portal Custom File


Select an Archive from the list and click on the Edit icon  to bring up the Customization Workflow Wizard. You cannot edit the File Name; however, you can edit the Archive File and apply this edited Archive File to the same switch(es) or a different switch(es). Follow the steps as described above to apply an edited Archive File.

Applying a Captive Portal File

You can also apply existing Archives to switches. Select an Archive File in the list, click on **Push to Switches** button, select the switches to which you want to apply the file, then click on the **Push File** button. Click **OK** to return to the Configuration Screen. The new Custom Archive that you apply will replace any existing Custom Archive on the switch(es).

Note: To "unassign" an Archive File from a switch, select it and click on the **Push to Switches** button. The Apply to Devices Screen appears. The switches to which the field is assigned will appear on the right. Move the switches from which you want to "unassign" the file to the left side of the screen and click the **Push File** button. Note that when a file is applied to a switch, it will override existing files. The switch will then present the default Captive Portal web pages to the user.

Deleting a Captive Portal Customization

Select a file and click on the Delete icon , then click **OK**, at the confirmation prompt. If the file has been assigned to a switch, a second prompt will appear. Click **OK** on the second prompt to delete the file.

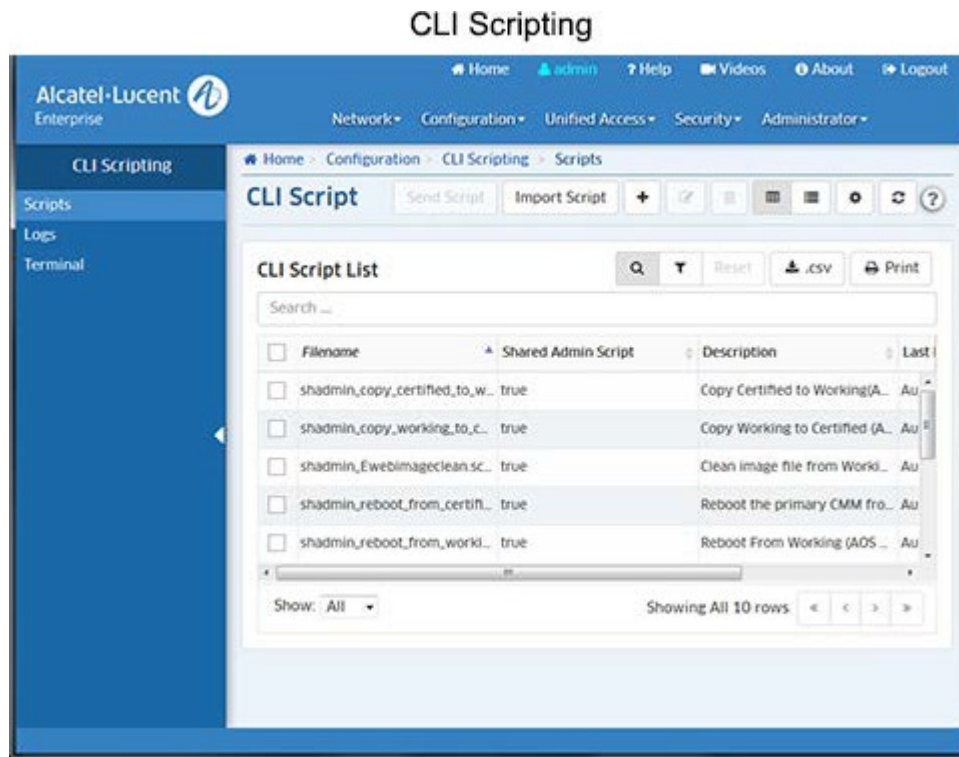


Captive Portal View

The Captive Portal View Screen is used to display Captive Portal configurations on specific switches. Click on the **Browse** button and select a switch. The Captive Portal configuration for the selected switch is displayed - Configuration, Profile, Profile Domain Policy List, and Domain Policy List. Click on each to view configuration details.

9.0 CLI Scripting

The CLI Scripting application is used to connect to and configure network devices via the CLI; and to configure multiple devices by creating CLI Script Files and applying the files to devices on the network.



The following screens are used for CLI Scripting:

- **Scripts** - Used to create CLI Scripting Files. You can then use these text-based files to configure one or more devices by applying the file via a Telnet or SSH session. You can also import existing text-based script files. Scripts can be created and stored on the local client, or can be saved to the OmniVista Server so they can be available to other clients ("shared").
- **Logs** - Displays log files of all CLI Scripts that have been applied to network devices. You can click on a log file to view CLI Scripting results on a command-by-command basis to see if the contents of the file were successfully applied to a device(s).
- **Terminal** - Used to establish CLI Scripting sessions by logging in to one or more devices and configuring the devices via supported CLI commands. Note that OmniVista CLI Scripting also supports SSHv2 enhanced encryption.

CLI Script

The CLI Scripting CLI Script Screen displays a list of all configured CLI Script Files. It is also used to create, import, edit, and delete, script files, and send script files to network devices. A CLI Script File is a text-based file used to configure one or more devices through OmniVista's CLI Scripting feature. CLI Scripting is especially useful in applying batch updates or common configurations across multiple devices. When a script file is applied, each command in the file is sent to the device. A user can create a CLI Script that is only available to that user, or can create a "shared" script that is available to any network administrator.

Note: Before attempting to send a script, OmniVista must know the CLI/FTP user name and password for each device being configured. If necessary, go to the Discovery application to specify CLI/FTP user name and password. You can also specify the CLI/FTP user name and password using the "Discovery - Edit Device" operation in the Topology application.

Pre-Configured Scripts

OmniVista includes pre-configured CLI Script Files, which are displayed in the CLI Script List (along with any user-configured scripts). These pre-configured scripts are "shared" scripts and are available to any network administrator. A brief description of each script, as well as the contents of the script file, are provided in the Details Panel for the script. Click on a script in the CLI Script List to view script details.

Important Note: Use caution when using the **shadmin_Ewebimageclean** script. Use the Resource Manager application to perform a full backup on the switch prior to an upgrade.

Creating a Script File

Click on the Create icon **+** . Enter a **Filename** for the script (e.g., show_switch). The file extension ".script" will be added automatically when the script file is saved. Select the **Shared Admin Script** checkbox to create a "shared" script that can be used by network administrators. If you do not select the checkbox, the script will only be available to you. When a "shared" script is created, the prefix "shadmin" is automatically assigned to the filename. To add a description of the script that will appear in the Details Panel, enter the description as follows at the top of the script. For example:

```
<js>
/* @@Enter a description here @@*/
```

Enter the commands to be applied to the switch via this script in the **Commands** field. Enter one command per line. The script can be a combination of both CLI commands and JavaScript. You can also define variables, or use OmniVista built-in variables to be used in the script. Note that if a script with the same name currently exists, an error message will appear. Re-name the script and continue as described above. Verify that the syntax of all the commands is correct, then click the **Add** button. The filename will be listed in the CLI Script List in alphabetical order. Select the script in the CLI Script List and click on the **Send Script** button to send the script to a network switch(es).

CLI and Java Scripts

Scripts can be a combination of both CLI commands and JavaScript. The following is an example of a CLI script containing JavaScript:

```
----- script start -----
<js>
var devtype = cli.deviceType();
if (devtype.indexOf("OS68") > -
1) cli.sendCmd("ls");
else if (devtype.indexOf("OS62") > -
1) cli.sendCmd("dir");
else
cli.sendCmd("files
");
if (devtype.indexOf("OS68") > -1)
{
cli.setTimeout(3, 30);
cli.sendCmd("show log
swlog");
}
else if (devtype.indexOf("OS66") > -1)
```

```

{
cli.setTimeout(5, 0);
cli.sendCmd("show log
swlog");
}
println("I got: " + cli.lastResponse() );
cli.sendCmd("ls " + "$USERVAR" ); /* user defined variable! */

</js>
----- script end -----

```

Notice in the above example, the usage of the variable 'cli' . This is a built-in variable that can be used within the scripting blocks. CLI offers the following functions:

- **sendCmd(String cmd)** allows the user to send a CLI command to the switch.
- **lastResponse()** returns a string that represents the switch output from the last command the user sent to the switch (whether the command was sent via JavaScript or just entered as CLI in the cli script itself). **deviceType()** returns the same string as can be seen via the Topology applications "Type" column.
- **setTimeout(minutes, seconds)** allows a caller to specify a hint to the CLI Scripting application about how long it could take for the next command to return a response. In the example above, the JavaScript specifies a timeout of 3 minutes and 30 seconds to apply to the next command (show log swlog) if the device is something like an OS6800-48. It specifies 5 minutes if the device is something like an OS6624. Some commands can be slow in returning output to the CLI Scripting/SSH session, so this can help prevent the scripting session from timing out before a response is received. Once the session is receiving a response from the command (e.g., show log swlog), the default timeout will be automatically reset. The user specified timeout does not take affect for the entire session, just the CLI command used after the call to **setTimeout(minutes, seconds)**. You may specify "0" for minutes or seconds according to what is needed. Negative numbers are converted to '0' internally, thus ignored.

If both minutes and seconds contain either "0" and/or negative numbers, the timeout request will be ignored. Therefore the minimum timeout will be 1 second (ex: `cli.setTimeout(0, 1);`).

- **trace(String message)** logs any arbitrary string passed as its 'message' argument to the CLI Scripting Audit Log. Can be contained in a variable for instance.
- **expectPrompt(String prompt)** sets-up the particular script (running on whatever devices) to expect a prompt that is not in the default collection of expected prompts. In other words, it allows the user to temporarily add to the set of prompts that CLI Scripting is hard-coded to recognize.
- **deviceType()** returns a string that contains the device's type as seen in the Topology application.
- **cliSleep(milliseconds)** allows the user to set a time, in milliseconds, before the next CLI command is executed.
- **errorLog(String message)** logs any 'error' argument to the CLI Scripting Audit Log. Can be contained in a variable for instance.
- **cli.sendCmd(more)** by default, the switch will stop running a command at a confirmation prompt and wait for the user to confirm the action. The 'more' command tells OmniVista to expect a specific prompts(s) and continue running the script. For example, the 'more' variable should be used when sending a script to reload the working directory (**reload working no rollback-timeout**). In the example below, the 'more' command tells OmniVista to expect the "Confirm Activate" and "Confirm New Activate" prompts, and to reload the switch from the Working Directory in 10 minutes.

```
----- script start -----
```

```

</js>
cli.sendCmd("more
");
cli.expectPrompt("Confirm Activate (Y/N) :");
cli.expectPrompt("Confirm New Activate (Y/N) :");
cli.sendCmd("reload working no rollback-timeout in
10:10"); cli.sendCmd("y");
</js>

```

----- script end -----

Enter only one command per line. Operational commands that automatically issue a confirmation prompt and require the user to type a response (such as, Y or N) are not supported in CLI script files. Examples include **takeover**, **reload**, **fsck**, etc. Do not attempt to include these command types in the script file. Instead, manually issue them via the standard CLI command line prompt. These operations can also be issued on a device-by-device basis via WebView or OmniVista.

Important Note: If a command that takes a long time to complete (e.g., “write memory flash-synchro”), is issued as the last command in a CLI script, the session can end right after the command is issued, ending the session before the command is executed. To avoid this problem, either add another command at the end of the script (“show chassis”), **or** add a taps timeout. For example, the following command sets a timeout of 0 minutes and 15 seconds:

```
<taps> set timeout 0 15 </taps>
```

Built-In Variables

OmniVista built-in variables are listed below.

- **\$BASE_MAC** - Replaced automatically with target base MAC address.
- **\$BOOT_DIR** - Replaced automatically with target boot directory (ex: working).
- **\$CHASSIS_TYPE** - Replaced automatically with target chassis type.
- **\$IP_ADDRESS** - Replaced automatically with target switch IP address.
- **\$LOGIN_ID** - Replaced automatically with target CLI/FTP User Name.
- **\$LOGIN_PWD** - Replaced automatically with target CLI/FTP Password.
- **¶\$READ_PWD** - Replaced automatically with target community string for SNMP reading.
- **\$SECOND_PWD** - Replaced automatically with the value of secondary password in the Discovery list item, if applicable. The secondary password for a device is set in the Edit Discovery Manager Entry window in the Discovery application.
- **\$SYS_LOCATION** - Replaced automatically with the location of the device as defined in sysLoction MIB-II variable.
- **\$SYS_NAME** - Replaced automatically with the name of the device as defined in sysName MIB-II variable.

- **\$\$SYSTEM_OID** - Replaced automatically with target unique object ID.
- **\$\$SYS_VERSION** - Replaced automatically with the MPM Version of the device as displayed in OmniVista.
- **\$WRITE_PWD** - Replaced automatically with target community string for SNMP writing.
- **cli.forgetPrompt()** - This CLI script directive is used to reverse the **cli.expectPrompt()** directive, providing a way to ignore prompts that interfere with script execution.

Important Note: If you are using Built-In variables **within a Java Script** , the variable must be contained within quotes (e.g. "**\$BASE_MAC**"). If you are using Built-In variables outside of a Java Script, the quotes are not required.

Script Directives

A tag, called <tapps> allows certain directives to the CLI Scripting application. <tapps> does not use a scripting engine. It is a set of supported commands that tell the CLI Scripting application how to handle certain actions. For example, a user may write the following CLI script that uses all of the supported <tapps> commands:

```
<tapps> set timeout 5
</tapps> qos apply
<tapps> import another.script </tapps>
<tapps> second password </tapps>
```

set timeout: The above script specifies a timeout for the *qos apply* command. It performs the same function as the previous Java Script example, but the user does not need to specify seconds. However the user must always specify minutes (the minutes can be "0" if the user wants to specify the timeout only in seconds).

Examples:

As shown above, to set a timeout of 5 minutes, only the *minutes* parameter is required:

```
<tapps> set timeout 5
</tapps> qos apply
```

To set a timeout of 15 seconds, you must first specify 0 minutes , then 15 seconds :

```
<tapps> set timeout 0 15
</tapps> qos apply
```

To set timeout of 5 minutes and 15 seconds, you would enter:

```
<tapps> set timeout 5 15
</tapps> qos apply
```

Note: The *set timeout* command only applies to the next command in the script (e.g., *qos apply*). Thereafter, the timeout reverts back to its default.

import script: The import script directive tells the CLI Scripting application to insert the commands from the specified script at that spot in the current script. This allows re-use of scripts by other scripts. In the example above, if the CLI Scripting application script named "another.script" contained only the command 'ls' , then 'ls' would be inserted at runtime at that point in the current script. The log output for a running of the current script would show the command 'qos apply' sent, followed by the command 'ls' being sent. Detection of loops takes place at strategic points in the CLI Scripting application on both the client and server sides.

second password: The second password directive tells the CLI Scripting application to prepare to send the password again. Some devices have a second login capability that requires the use of a second password. This second password for a given device is set by the user via Topology when a device is selected for Editing. The value in the Topology 'Secondary Password:' field will be used by this new <tapps> feature as the password to set when or if the device prompts for a password.

last command: On some devices (e.g., OA5510-TE), commands such as 'reload' will 'hang' the OmniVista CLI Scripting session because the switch session will end without closing the session with OmniVista. The 'last command' directive, <tapps> lastcmd </tapps>, alerts OmniVista that the next command is the last command and a response may not be received after this command. OmniVista will gather whatever response is given before reload and close the session. For example:

```
<js>
cli.sendCmd("enable
");
cli.sendCmd("$SECOND_PWD");
cli.expectPrompt("Do you want to save config before rebooting
(y/[n])"); cli.expectPrompt("Do you really want to reboot the Chassis
(y/[n])");
</js
>
relo
ad n
<tapps> lastcmd
</tapps> y
```

The expectPrompt() calls in the java scripts train the CLI Script to send the next value upon receiving the specific prompt from the switch. Note that 'lastcmd' is used before "Y" and not reload command.

Importing a Script File


Although OmniVista allows users to manually create script files within the CLI Scripting application, existing script files can also be imported. In other words, a file containing a set of CLI commands can be accessed from a server or local drive and then applied to one or more devices. This allows users to maintain a library of network configurations and then apply them to devices in their network as needed. Before importing a file to one or more devices, consider the following important guidelines:

- Any script file being imported must be text-based (ASCII).
- Although file extensions such as .txt and .ascii are supported, the file extension .script is recommended.
- All CLI commands contained in the file must be supported on the device. Also, operational commands that automatically issue a confirmation prompt and require the user to type a response (such as, Y or N) are not supported in CLI script files. Examples include **takeover** , **reload** , **fsck** ,
- etc. CLI commands must also be entered into the text file *one command per line* .
- Only one script file can be imported at a time.

To import a script file, click the **Import Script** button at the top of the screen. On the **Import Script** window, click on the **Browse** button to locate the file. Select the file and click on the **Import** button, then click **Finish**. The script will be imported as a "shared" script with the current date appended to the script name (e.g., new_script20161026.script).

Note: If you are browsing for a file with an extension other than **.script** , be sure to select **Files of Type -> All Files** in the dialog box.

Editing a Script File

Select the script in the CLI Script List and click on the Edit icon . Edit the script commands and click on the **Apply** button. Note that you cannot edit the script name or "shared" status. If you want to change the "shared" status of script, delete the script and re-create it.

Important Note: When the changes are saved, the previous contents of the script file are overwritten. To preserve the original contents of the file, be sure to make a backup copy before editing.

Sending a Script File

You can send a script file to a single device or multiple devices in the network. Select a script in the CLI Script List and click on the **Send Script** button to bring up the Send Script Wizard. Complete the screens as described below.

Script Info

The name and contents of the selected script file are displayed. Click **Next**. (Note that you have the option of selecting a different script. Click on the **Browse** button to bring up all of the scripts in the CLI Script List, select a script and click **OK**, then click **Next**.)

Device Selection

Select an option from the drop-down menu (User Switch Picker/Use Topology) and click on the **Add/Remove Devices** button to select devices..

- **Use Switch Picker** - Select the devices and click **OK**. Click **Next**.
- **Use Topology** - The Topology application will launch in the Physical Map view. Select the device(s), then click on the **OK** button at the bottom of the Detail Panel to return to the Send Script Wizard. The devices will appear in the list of devices. Click **Next**.

Scheduler


You can send a script immediately to the selected device(s), or schedule the script to be run at a specific time or at regular intervals. After selecting/configuring an option, click **Next**.

- **Now** - The script will be sent immediately to the selected device(s)
- **Periodically** - Schedule the script as described below.
 - **Start Time** - Set a time to start the repeating script.
 - **End Time** - Set a time to stop the repeating script.
 - **Simple** - Select this option to configure a repeating script. Set an **Interval** using the Days/Hours/Minutes/Seconds fields, and enable the **Repeat** field to the number of times to repeat the script until the configured "End Time" is reached.
 - **Cron** - Select this option and use the drop down menus in each tab (e.g., Minute, Hour, Day) to configure a repeating cron job. The cron job will continue until the configured "End Time" is reached.

Define User Variables

If there are variables within the script, the **Define User Variables** Screen is displayed. Click in field next to the variable and enter value to be used. After completing all of the variable fields, click the **Send Script** button at the bottom of the screen.

Deleting a Script File

Select the Script File in the CLI Script List and click on the Delete icon . Click **OK** at the confirmation prompt. Note that when a file is deleted, it is permanently removed from the scripting_files directory, and cannot be recovered.

CLI Script Details

The CLI Script List displays basic information about all configured CLI Scripts stored on the OmniVista Server. Click on a script to view the commands contained in the script.


- **File Name** - The script Filename.
- **Shared Admin Script** - Whether or not the script is a shared script (True) or not (False).
- **Description** - A brief description of the script.
- **Commands** - The commands contained in the script.

Logs

The CLI Scripting Logs Screen displays a list log files of all CLI Scripts that have been applied to network devices. You can click on a log file to view CLI Scripting results on a command-by-command basis. In other words, it displays whether the contents of a file were successfully applied to the device. A log file also provides a record of a particular configuration, as well as effective troubleshooting information, when applicable. The screen can be used to view, export, or delete CLI Scripting Logs.

Note: As with the scripting files, log files are automatically stored on the OmniVista Server or local system. File locations may vary, depending on the OmniVista installation, but can generally be found at a path similar to the following: Alcatel OmniVista 2500\data\cli_scripting_logs. By default, log files are placed in a directory indicating the IP address of the corresponding device.


Displaying a Log File

Click on a log file to display the contents. You can look through the file to view the application of the Log File on a command-by-command basis. Unless an error has occurred, the log file will closely resemble the script file (i.e., it will list only the CLI commands that were applied to the device). If an error occurs, an error notification is displayed in the log, following the CLI command that triggered the error. You can search for a command or specific text string in the log file by entering the text in the Search field and clicking on the **Search Next** or **Search Previous** buttons. You can also configure a filter to view specific information by clicking on the Filter icon  and creating/selecting a filter.

Exporting a Log File

You can export a log file to another location (directory). Select the file and click on the **Export** button. Browse to the location to which you want to export the file and click **OK**.

Deleting a Log File

Select the log file(s) and click on the Delete icon . Click **OK** at the confirmation prompt.

Terminal

The CLI Scripting Terminal Screen is used to establish a basic CLI Scripting session with a device. You can locate and connect to a device using a Switch Picker or the Topology Map. Once you are connected to a device, log into the device to issue CLI Scripting commands. Note that OmniVista must know the CLI/FTP user name and password for a device to log into the device. If necessary, go to the Discovery application to specify CLI/FTP user name and password. You can also specify the CLI/FTP user name and password using the "Discovery - Edit Device" operation in the Topology application.

Connecting to a Device

You can locate and connect to a device using a Switch Picker or the Topology Map.

Switch Picker

To connect to a device using the Switch Picker, select **Using Switch Picker** from the drop-down menu and click on the **Browse** button. Select the switch you want to connect to and click **OK**. Click on the **Telnet** button or the **SSH** button at the top of the screen to connect to the device. Log into the device to begin the session. To disconnect from a device, click on the **Disconnect** button at the top of the screen, or enter *exit* at the command prompt.

Topology Map

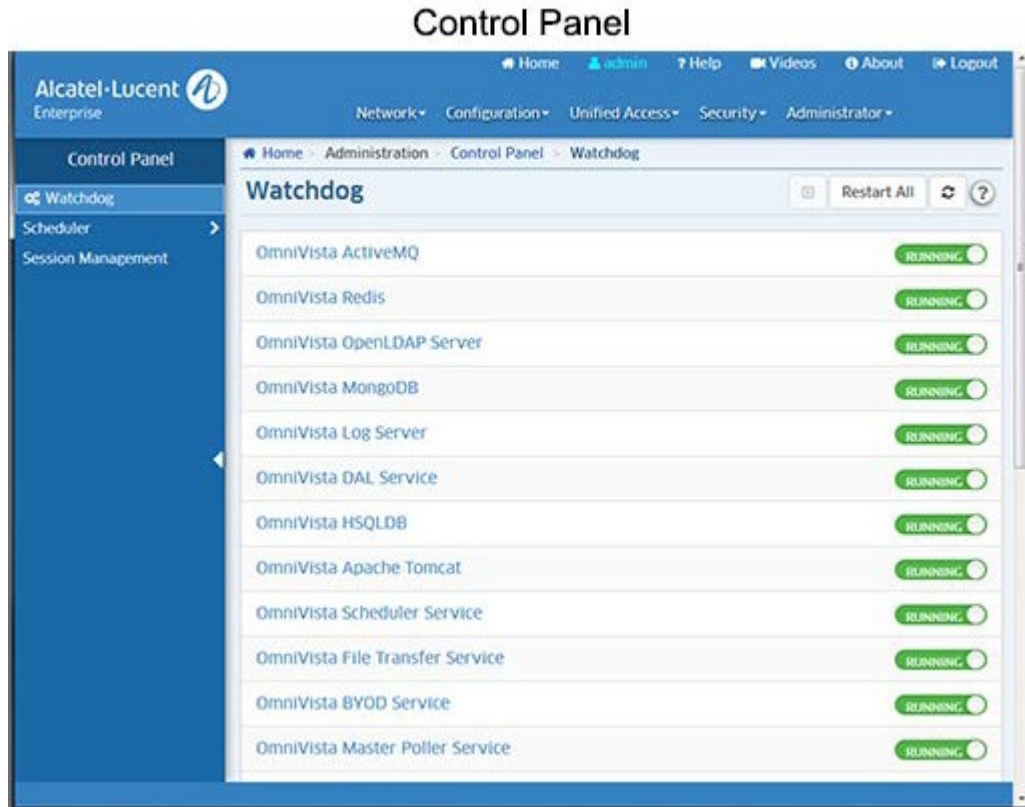
To connect to a device using the Topology Map, select **Using Topology** from the drop down menu and click on the **Browse** button. The Topology application will open. Select the device you want to connect to and click on the **Telnet** link in the Selection Operations area on the right side of the screen. Log into the device to begin the session. To disconnect from a device, enter *exit* at the command prompt.

Session Preferences


SSH (Secure Shell) provides CLI Scripting sessions with enhanced encryption and security. SSH may be mandatory for some device types. OmniVista uses SSH by default for those devices requiring SSH. However, for AOS and other devices where SSH is optional, standard Telnet is the default setting. To use SSH, you must specify SSH either on a device-by-device basis, or on multiple devices using the Discovery Profile feature in the Discovery application.

10.0 Control Panel

The Control Panel application is used to access the Watchdog, Scheduler, and Session Management features. The Watchdog feature displays the status of all of the services used by OmniVista; and is used to start/stop services. The Scheduler feature provides an overview of all currently Scheduled jobs (System Jobs and User-Defined Jobs), and is used to start/stop, edit, or delete a User-Defined Job. The Scheduler feature also provides a history of all completed Scheduler jobs. The Session Management feature displays a list of all OmniVista Client login sessions, and can be used to log out a session.



Watchdog

The Control Panel Watchdog Screen displays the status of all of the services used by OmniVista (Running/Stopped). Click on any service to view an information panel for the service (e.g., description, status, dependencies, statistics). To start/stop a service, click on the slider control next to the service (**Running/Stopped**). If you are stopping a service, click **Yes** at the confirmation prompt. Click on the Start All icon  to start all stopped services. To stop and restart all services, Click on the **Restart All** button.

You can also stop/start a service in the information panel. To stop a service, click on the **Stop** button, then click **Yes** at the confirmation prompt. This will stop the service and all of its dependent services. To start a Service, click the **Start** button to start the service and all of its dependent services.





Warning: If you stop certain services (e.g., ActiveMQ, Apache Tomcat) or a service that these services depend on, the web server will shut down, and you will have to restart the service manually. You will receive a warning prompt whenever you try to shut down one of these services.

Scheduler Jobs

The Control Panel Scheduler Jobs Screen provides an overview of all currently Scheduled jobs (System Jobs and User-Defined Jobs). System Jobs are automatically scheduled by OmniVista. System Jobs cannot be edited or deleted. User-Defined Jobs are scheduled by users within OmniVista applications (e.g., using the Resource Manager application to scheduled backup job). To view specific details about a job, click on the job in the table to display job details (e.g., Start Time, End Time, Cron Description). You can also start/stop, edit, or delete a User-Defined Job. Note that you can only view System Jobs. You cannot start/stop, edit, or delete these jobs.

Starting/Stopping Scheduled Jobs

You can start/stop/pause a job by selecting the job and clicking on the applicable icon:

- **Start**  - Register the job in the schedule and start executing immediately if its start time is in the past.
- **Stop**  - Stop the current job. Stopped job will execute normally in the next cycle.
- **Pause**  - Stop the current execution and save its progress state. The job also is removed from the schedule and will not be executed at next trigger. You can restart the job by selecting it and clicking on the Start icon . The job will be started job from the last state and resume the job schedule.

Viewing Scheduler Jobs

The Scheduler Jobs Table lists all schedule System and User-Defined Jobs. Click on the applicable tab at the top of the table to view a list of each type. The table provides the basic information. Click on a job to view detailed information.

Basic Information




- **Name** - The system-generated job name.
- **Group** - The system-generated job group. A job group is a logical grouping of related jobs grouped by application, framework, etc. (e.g., Analytic, Poller, VM Snooping). You can sort or search on a job group in the Scheduler Jobs Table to view related jobs.
- **Status** - The status of the job (e.g., Scheduled, Waiting).

Detailed Information

- **Name** - The system-generated job name.
- **Group** - The system-generated job group. A job group is a logical grouping of related jobs grouped by application, framework, etc. (e.g., Analytic, Poller, VM Snooping). You can sort or search on a job group in the Scheduler Jobs Table to view related jobs.
- **Priority** - The job priority. If jobs are initialized at the same time, the job with the higher priority will begin first (Range = 1 - 10).
- **Actor** - The system-generated behavior description for the job.
- **Select Device Type** - The method for selecting devices, if applicable (e.g., Specific Devices)
- **Device List** - A list of devices based on the Selected Device Type if applicable.

- **Overlap Policy** - The Overlap Policy determines the action OmniVista will take if there is a job overlap
 - **Ignore When Overlap** - The next run (cycle) of the job will be skipped if it is still being executed at the scheduled time.
 - **Replace When Overlap** - The job will start fresh (restart) in the next run (cycle) if it is still being executed at the scheduled time.
- **Action From Crash Policy** - The Crash Policy determines the action to take if the job crashes before completion:
 - **Start Afresh From Crash** - The job will start fresh in the next run (cycle) if it is in a failed state at the scheduled time.
 - **Resume From Crash** - The job will resume from the failure point in the next run (cycle) if it is in a failed state at the scheduled time.
- **Start Time** - The configured start time for the job.
- **End Time** - The configured end time for the job.
- **Schedule** - The schedule type for the job:
 - **Simple** - The job repeats at specific intervals.
 - Interval - The repeat interval for the job (e.g., 1 Day, 1 Hour). Repeat - The number of times the job will repeat.
 - Retry Count - The number of times the job will retry after a failure. Retry Interval - The duration from failure to next retry, in seconds.
 - Timeout - The maximum amount of time the job will run before timing out, in seconds. Owner - The user who created the job (e.g., admin). User-Defined Jobs only.
 - **Cron** - The job is a recurring Cron job.
 - Cron Description - A brief description of the Cron Job.
 - Retry Count - The number of times the job will retry after a failure. Retry Interval - The duration from failure to next retry, in seconds.
 - Timeout - The maximum amount of time the job will run before timing out, in seconds. Owner - The user who created the job (e.g., admin). User-Defined Jobs only.


Editing a Scheduled Job

You must be an admin user to edit a scheduled job. To edit a job, select the job in the Scheduler Jobs table and click on the Edit icon . Edit the fields as described below and click on the **Save** button. Note that you can only edit a "Paused" or "Waiting" Scheduler job. If necessary, select the job you want to edit and click on the Pause icon . When you are done editing the job, click on the Start icon  to activate the job.


- **Name** - The system-generated job name. This field cannot be modified.
- **Group** - The system-generated job group. A job group is a logical grouping of related jobs grouped by application, framework, etc. (e.g., Analytic, Poller, VM Snooping). You can sort or search on a job group in the Scheduler Jobs Table to view related jobs. This field cannot be modified.
- **Priority** - The job priority. If jobs are initialized at the same time, the job with the higher priority will begin first (Range = 1 - 10).
- **Actor** - The system-generated behavior description for the job. This field cannot be modified.
- **Overlap Policy** - Sets the Overlap Policy that determines the action to take if there is a job overlap:
 - **Ignore When Overlap** - The next run (cycle) of the job will be skipped if it is still being executed at the scheduled time.
 - **Replace When Overlap** - The job will start fresh (restart) in the next run (cycle) if it is still being executed at the scheduled time.

- **Action From Crash Policy** - Sets the Crash Policy that determines the action to take if the job crashes before completion:
 - **Start Afresh From Crash** - The job will start fresh in the next run (cycle) if it is in a failed state at the scheduled time.
 - **Resume From Crash** - The job will resume from the failure point in the next run (cycle) if it is in a failed state at the scheduled time.
- **Start Time** - Enable this option and schedule a specific start day and time for the job. You can enter the date and time in the field or use the drop-down calendar to select the day, and then edit the time in the field. Note that if the start time is before the current time, the job will start immediately.
- **End Time** - Enable option field and schedule a specific end day and time for the job. You can enter the date and time in the field or use the drop-down calendar to select the day, and then edit the time in the field.
- **Schedule**
 - **Simple** - Select this radio button and schedule the job to repeat at specific intervals (e.g., Days, Hours, Minutes, Seconds). Enable the **Repeat** option to limit the number of times the interval will be repeated.
 - **Cron** - Select this radio button to schedule the job as a recurring Cron job.
 - **Retry** - Enable this option and configure the job retry option: Count = how many times the job will retry after a failure. Interval = the duration from failure to next retry, in seconds. (Count Range = 0 - 99, Interval Range = 0 - 99)
 - **Timeout** - Enable this option and configure the maximum amount of time a job will run before timing out, in seconds (Range = 20 - 9999). If it is disabled, a job execution could run forever.

Deleting a Scheduled Job

To delete a job, select the job in the Scheduler Jobs table and click on the Delete icon . Click **OK** at the confirmation prompt. The job will be deleted and will no longer run.

Scheduler History

The Control Panel Scheduler History Screen displays a historical overview of all completed Scheduler jobs (e.g., device audit, license audit). Click on a job to view specific details about the job. You can manually remove an event(s) by selecting the event(s) and clicking on the Delete icon . Note that you must be an "admin" user to view the Scheduler History Screen.

Session Management

The Control Panel Session Management Screen displays a list of all OmniVista Client login sessions, and can be used to log out a session. A single user can have multiple sessions, logging into the server from different clients. Logging out one session will not affect other sessions of same user. To log out of a session(s), select the session(s), click on the **Logout Selected Sessions** button, and click **Yes** at the confirmation prompt.

Session Information

The Session Management Table displays the following information about each user who logged into the server:

- **User Name** - The user name.
- **Host Name** - The host name of the client where this session originated. If the address cannot be resolved to a host name, the IP address is displayed.
- **IP Address** - The IP address of the client where this session originated.

- **First Name** - The first name of the user.
- **Last Name** - The last name of the user.
- **Description** - The user description.
- **Login Server** - The name of the Authentication Server used to authenticate the session.
- **Logged in Since** - A timestamp of when was the session was created.
- **User Groups** - The user group(s) to which the user belongs.

11.0 Discovery

The Discovery application is used to discover network devices. The information gathered is used by OmniVista applications to view and configure the network. The information includes:

- Alcatel-Lucent Enterprise devices in the network.
- The links between devices in the network. This information is used to display network links in graphical maps of network regions.
- Additional link information required by OmniVista's Locator application. Third-party devices built by Cisco and Extreme.
- Any additional third-party devices for which support has been added.

Discovery

The screenshot shows the Alcatel-Lucent Enterprise Discovery application interface. The main content area is titled "Managed Devices" and contains an "Inventory" table. The table has the following columns: Friendly Name, Name, Address, Status, DNS Name, Type, and Version. The status for all devices is "Up".

Friendly Name	Name	Address	Status	DNS Name	Type	Version
10.255.225.97	OAW-4030	10.255.225.97	Up		OAW-4030	6.4.4.9.55980
10.255.225.29	NMS-IAP	10.255.225.29	Up		OAW-AP115	6.4.4.6-4.2.4.0.5
10.255.225.96	OAW-4030	10.255.225.96	Up		OAW-4030	6.4.4.9.55980
10.255.225.99	Aruba-4704	10.255.225.99	Up		OAW-4704	6.4.4.9.55980
10.255.221.30	NMS-TELDAT-30-test	10.255.221.30	Up		OA5710V	11.00.00.02.05
10.255.225.42	name	10.255.225.42	Up		OS6350-48	6.7.1.72.R04
10.255.225.153	switch153	10.255.225.153	Up		OS6250-24	6.7.1.50.R02
10.255.225.201	name	10.255.225.201	Up		OS6400-24	6.4.5.635.R02
10.255.225.62	name	10.255.225.62	Up		OS6450-C24	6.6.5.38.R02
10.255.225.230	name	10.255.225.230	Up		OS6865-P16X	8.3.1.309.R01
10.255.221.149	name	10.255.221.149	Up		OS6200-12LS	1.7.1.10

Discovery is configured and performed using the screens below:

- **Managed Devices** - Displays a list of all discovered network devices. It is also used to discover/re-discover devices.
- **Discovery Profile** - Used to create a profile containing the parameters used for discovery (e.g., SNMP version, permissions).
- **Third-Party Devices Support** - Used to configure discovery parameters for third-party devices (e.g., OID, Display Name).
- **Import MIBs** - Used to import new or updated MIB files into OmniVista.
- **Hardware Inventory** - Used to view inventory information (e.g., CMM, Chassis, Power Supplies) for any discovered device.
- **Ports** - Used to display information about ports on network devices, and is also used to enable/disable device ports.

- **Link** - Displays all links that were learned during the discovery process or created manually in OmniVista. It is also used to manually create, edit, and delete manual links.
- **Settings** - Used to configure automatic discovery frequency parameters, enable/disable IP Failover, and configure switch monitoring.

Managed Devices

The Discovery application Managed Devices Screen displays a list of all discovered network devices (default). It is also used to discover/re-discover devices and add, clone, edit, delete, and search for devices. You can also perform certain operations on devices such as ping/poll devices, configure traps, locate end stations, and reboot devices.

Note: Admin and Netadmin users will see all discovered network devices. For other users the devices displayed depend on the User Role and User Group as defined in the Users and User Groups application. Only the devices in the maps associated with a User's Role will be displayed.

Discovering/Re-Discovering Devices

You can discover new network devices or re-discover devices to update information for those devices.


Discovering Devices


OmniVista performs a discovery based on a specified IP address range and Discovery Profile. The Range specifies the IP address range in which you want to discover devices. The Range is associated with a Discovery Profile. The Discovery Profile contains the parameters that are used by OmniVista when performing the discovery (e.g., SNMP version used to discover devices, FTP/Telnet passwords needed to connect to a device).

Discovering Devices Using and Existing Range

To discover devices using an existing range, click on the **Discover New Devices** button at the top of the screen. Any configured ranges appear in the Ranges List. Select a range and click on the **Discover Now** button. When you click on the **Discover Now** button, the discovery will begin and a progress screen will appear. When the discovery is complete the discovered devices will appear on the Discovery Screen in the Inventory List.

Creating a New Range for Discovering Devices

If you want to create a new Range, click on the **Discover New Devices** button to bring up the Ranges List. Click on the Create icon  and complete the fields as described below. When you are finished, select the new Range in the Ranges List and click on the **Discover Now** button. When you click on the **Discover Now** button, the discovery will begin and a progress screen will appear. When the discovery is complete the discovered devices will appear on the Discovery Screen in the Inventory List.

- **Start IP** - The starting IP address of the discovery range (e.g., 10.255.10.1)
- **End IP** - The ending IP address of the discovery range (e.g., 10.255.10.254)
- **Subnet Mask** - The subnet mask used for the discovery range (e.g., 255.255.255.0)
- **Description** - A description for the discovery range.
- **Choose Discovery Profiles** - Select the Discovery Profile(s) to use for the discovery. If necessary, click on the Add icon  to go to the Discovery Profiles Screen and create a new profile. After the profile is created, return to this screen and create a new range using the new profile(s). You **must**

associate a discovery range with at least one Discovery Profile to perform the discovery.

If you want to use more than one Discovery Profile for new range, drag and drop the profiles in the list to prioritize the order in which they are used for discovery. OmniVista will first attempt to discover each device using the first profile listed. If OmniVista cannot communicate with a device using the first profile, it will try the next profile, and so on. If OmniVista cannot communicate with a device using any of the profiles, it will attempt to communicate with the switch using the default profile. Once OmniVista successfully communicates with and discovers a device, it will use that Discovery Profile for all future communications with that device (unless you edit the device in the Inventory List).

Re-Discovering Devices

You can "re-discover" previously-discovered devices to update information about a device(s). For example, you might wish to re-discover a device to learn VLAN information that was not gathered during the first discovery; or re-discover a device if that device was re-configured outside of OmniVista. To re-discover a device(s), select the device(s) in the Inventory List and click on the **Rediscover** button. The discovery will begin and a progress screen will appear. When the discovery is complete the discovered devices will appear on the Discovery Screen in the Inventory List.

Note: If a switch was discovered previously, and a new discovery is performed using a Discovery Profile that specifies different parameters (e.g., CLI/FTP user name and password, Shell Preference), the profile values will not overwrite the values already specified to OmniVista for that device. The values specified in the profile will apply to newly-discovered devices only.

Adding a Device

You can manually add a device to the Inventory List. Click on the Create icon **+** bring up the Add New Device Screen and complete the fields as described below. After completing the fields, click on the **Add** button to add the device.

Note that you can "clone" an existing device to quickly add a new device. Select a device in the Inventory Table and click on the **Clone** button at the top of the screen. The Add New Device Screen will appear with all of the fields reflecting the configuration of the selected device. Enter the IP address for the new device and, if necessary, edit any fields. After completing the fields, click on the **Add** button to add the device.

General

- **Device Name** - The user-configured device name (display only)
- **IP Address** - The device's primary IP address.
- **CLI/FTP User Name** - The user name that OmniVista will use to establish CLI/FTP sessions with the discovered devices. The user name specified will be used to auto-login to devices when CLI sessions are established. It will also be used to perform FTP with the device when configuration files are saved and restored (see note below).
- **CLI/FTP Password** - The password that OmniVista will use to establish CLI/FTP sessions with the discovered devices. The user name specified will be used to auto-login to devices when CLI sessions are established. It will also be used to perform FTP with the device when configuration files are saved and restored (see note below).
- **Confirm CLI/FTP Password** - Confirm the CLI/FTP Password.
- **Secondary Password** - The secondary CLI/FTP Password, if applicable.
- **Confirm Secondary Password** - Confirm the secondary CLI/FTP Password, if applicable.

Note: The CLI/FTP User Name and Password fields enable you to inform OmniVista of the device's CLI/FTP User Name and Password. A device's CLI/FTP User Name and Password cannot be configured from OmniVista, they must be configured directly on the device. If you do not define the CLI/FTP User Name and Password and you attempt to save, restore, or upgrade configuration files for a device, you will be individually prompted for the CLI/FTP User Name and Password of each individual device for which configuration files are being saved, restored, or upgraded. Also, OmniVista will be unable to auto-login to the device when establishing CLI Scripting sessions.

SNMP

- **SNMP Version** - The SNMP version that OmniVista will use to communicate with the device. The default version for AOS devices is v2, but v1 and v3 are also supported.
- **Timeout (msec)** - The time period, in milliseconds, that OmniVista will wait for a switch to respond to a connection request before assuming that the request has timed-out (Default = 5,000).
- **Read Community** - The device's "get" community name. The "get" community name enables OmniVista to read information from the device (see note below).
- **Write Community** - The device's "set" community name. The "set" community name enables OmniVista to write information to the device (see note below).
- **Retry Count** - The number of times that OmniVista will attempt to connect to a switch (Default = 3).
- **User Name (SNMP v3 Only)** - The SNMP version 3 user name.
- **Auth & Priv Protocol (SNMP v3 Only)** - The authentication protocol OmniVista will use for SNMP communication with the device. Authentication uses a secret key to produce a "fingerprint" of the message. The fingerprint is included within the message. The device that receives the message uses the same secret key to validate that the fingerprint is correct. If it is, and if the message was received in a timely manner, then the message is considered authenticated. Otherwise, the message is discarded. The fingerprint is called a Message Authentication Code, or MAC. Note that if you are using SHA256+AES 192 or 256 authentication protocols you must download and install the Zulu Cryptography Extension Kit (CEK) using the Preferences application (Preferences - System Settings - Install Zulu CEK).
- **Auth Password (MD5/SHA)** - The password that OmniVista will use for the MD5 or SHA authentication protocol. This must be the same password that is defined on the switch for MD5 or SHA. If no authentication password is entered, neither authentication nor privacy encryption will be used.
- **Confirm Auth Password (MD5/SHA)** - Re-enter the Auth password.
- **Private Password (MD5/SHA)** - The password that will be used as the secret key. This must be the same password that is defined on the switch for the CBC-DES Symmetric Encryption Protocol. If an authentication password is entered, but no privacy password is entered, authentication will be used without privacy encryption.
- **Confirm Private Password (MD5/SHA)** - Re-enter private password.
- **Context Name (SNMP v3 Only)** - A unique context name for this context. An SNMP context is a collection of management information accessible by an SNMP entity, in this case OmniVista. A context identifies a subset of management information, in this case the management information OmniVista has about the individual device. OmniVista, as an SNMP entity, has access to many SNMP contexts: one for each device it manages. Each context must be identified by a unique context name and a unique context ID. Note that an item of management information may exist in more than one context.

Technically, the Context Name and Context ID provide a means of distinguishing specific instances of information in the MIB modules from the set of all instances of that information within the management domain.

- **Context ID (SNMP v3 Only)** - A unique context ID for this context. As explained above, each context must be identified by a unique context name and a unique context ID. Note that neither the **Context Name** nor the **Context ID** are required for AOS or default third-party devices supported by OmniVista. Leave these fields blank unless you are using a non-default third-party device that requires definition of a Context Name and Context ID.

Note: If a device's "get" and "set" community names are "public" (the default) you can leave these fields blank (OmniVista uses the default name (public) when the field is blank. The community names are not configurable from OmniVista, they must be configured directly on the device. Also note that when you use SNMP Version 3, community names are ignored.

Advanced Settings

- **Trap Station User Name** - The user name that will be used when an AOS device is configured to send traps to OmniVista. AOS devices require that a valid device user name be specified with the trap station configuration entry. If this field is left blank, the following switch user names will be used by default for trap station configuration entries:
 - If OmniVista is configured to use SNMP version 3 with this device, the SNMP version 3 user name entered for the device will be used as the switch user name in the trap station configuration entry.
 - If OmniVista is configured to use SNMP version 1 or SNMP version 2 with this device, the read community string for the device will be used as the switch user name in the trap station configuration entry.

When using SNMP version 1 or 2, switch user names are interchangeable with community strings as long as community string mapping is not in use on the switch. If community string mapping is not in use, and an AOS switch is discovered using SNMP version 1 or 2 with a default read community string of "public", or even with a non-default read community string such as "thomas", these community strings are valid switch user names for trap station configuration entries. In this case, no further configuration is required and this field can be left blank. However, if community string mapping is enabled on the device, the community string with which the switch is discovered is not guaranteed to be a valid device user name, and thus is not guaranteed to be a valid device user name for a trap station configuration entry. In this case, you should enter a valid device user name in the Trap Station User Name field.

- **Discover Link** - Specifies how OmniVista will discover the device's links to other devices (Normal, As OEM Devices). Select **As OEM Devices** to enable OmniVista to automatically discovery links using functionality from OmniVista's Locator application. This option is useful if you want to discover links on devices that do not support adjacency protocols. If a device does not support an adjacency protocol that enables OmniVista to discover physical links, the endstation search algorithms used by the Locator application are invoked at each polling cycle to discover the device's links. All links discovered are automatically displayed in Topology maps.

This approach works well for devices located at the edge of the network that do not support adjacency protocols. However, when a series of such switches are interconnected at the core of a network, this approach may "discover" more links than are meaningful. As an example, consider a series of such devices connected in a chain. Use of the Locator endstation search algorithms, without benefit of any actual knowledge of how the switches are connected, will result in showing links between all the devices as a "cloud" instead of a chain. Such situations can be corrected by adding explicit manual links. For example, in the situation described, adding manual links for the actual connections will solve the problem by giving OmniVista the knowledge it needs to show the


connections accurately.


- **Shell Preference** - OmniVista's CLI Scripting application supports both the Telnet and SSH command line interfaces. SSH (Secure Shell) is a Telnet-like utility that provides encryption and is far more secure than Telnet. If you select **SSH**, SSH will be used as the default command line interface for the device.
In addition, Secure Shell FTP will be used as the default FTP method in Resource Manager. If you select **Telnet**, Telnet will be used as the default command line interface for the device and regular FTP will be used as the default FTP method in Resource Manager. Ensure that devices are capable of SSH before you enable the **Prefer SSH** checkbox. OmniVista does not verify devices' SSH capabilities. All AOS devices are SSH-capable. (Default = SSH)
- **Use Get Bulk** - Enables/Disables Get Bulk Operations. The SNMP version 2 Get Bulk operation is used for retrieving large amounts of data, particularly from large tables. The Get Bulk operation performs continuous Get Next operations, each time requesting the number of table rows specified by the value in the **Max Repetitions** field (below). For example, if the value in the Max Repetitions field is 10, each Get Next operation will request 10 rows of table data. Note that the number of rows of data actually returned by the device will be determined by the amount of memory the device has available at that time. (Default = Enabled)
- **Max Repetitions** - The number of rows of table data that the Get Bulk operation will request in each Get Next operation, if enabled.

Cloning a Device

You can "clone" an existing device to quickly add a new device. Select a device in the Inventory List and click on the **Clone** button at the top of the screen. The Add New Device Screen will appear with all of the fields reflecting the configuration of the selected device. Enter the IP address for the new device and, if necessary, edit any fields. After completing the fields, click on the **Add** button to add the device.

Editing a Device


You can edit a single device or edit multiple devices at the same time. To edit a single device, select a device in the Inventory List and click on the Edit icon . The Edit Discovery Manager Entry Screen will appear. Edit any fields as described above and click on the **Apply** button. Note that you cannot edit the Device Name.

To edit multiple devices, select the devices in the Inventory List and click on the Edit icon . The Edit Discovery Manager Entry Screen will appear and the selected devices will be listed in the IP Address Field. You can click on the device list to bring up the Selected Devices Window to display device information. Click again anywhere on the screen to close the window. Edit any fields as described above and click on the **Apply** button.

Note that if a field has different values among the selected devices (e.g., different location, different password), the field will be blank (drop-down menu) or grayed out ("yes/no" slider field), and a "Click to Overwrite" link will appear beneath the field. If you want to configure a common value for the field on all selected devices, click on the "Click to Overwrite" link and enter a value. If you want the values on each device to remain as they are, leave the field as it is. If you change your mind and want the fields to retain their original values, click on the "Retain Original Values" link **before** clicking on the **Apply** button.

Note: When you edit a device, it is important to understand that you are editing OmniVista's knowledge of the device, not the device itself.


Deleting a Device

Select a device(s) and click on the Delete icon , then click **OK** at the confirmation prompt.

Searching for a Device

You can search for a device by keyword by entering the search criteria in the Search fields at the top of the Inventory List. Enter any search criteria based on the contents found in the table and the list will change to display only those device containing the search criteria.

Perform Device Operations


You can also perform certain operations on devices in the Inventory List such as ping/poll devices, configure traps, locate end stations, and reboot devices. Select a device(s) in the Inventory List. Click on the Operations icon  at the top of the Inventory List and select an option from the drop-down list. Note that not all operations are supported on all devices; and some operations can only be performed on a single device, not multiple devices. If an operation is not supported for the selected device(s), it will be grayed out in the list.

- **Ping** - Immediately pings the selected device(s). Progress is shown on the Progress Screen. Click the **Finish** button to return to the Discovery Screen.
- **Poll For Traps** - Immediately polls the selected device(s) for traps. A message is displayed at the top of the Discovery Screen when polling is complete. Traps can be viewed in the Notifications application.
- **Poll Links** - Immediately polls links on the selected device(s). Progress is shown on the Progress Screen. Click the **Finish** button to return to the Discovery Screen.
- **Configure Health Thresholds** - Brings up the Configuring Devices Health Thresholds Screen in the Discovery application. Health Thresholds are used to set limits for health traps. If a device has been configured to send health traps, a trap will be sent whenever a monitored item's current utilization exceeds the configured health threshold. Configure the CPU, Memory, or Temperature Threshold for the selected device(s) and click on the **Apply** button. Note that you cannot configure the Temperature Threshold. The Temperature Threshold is hard coded on devices. Also note that changes made to health thresholds will not appear until the next polling cycle (up to an hour).
- **Locator - Locate End Stations** - Launches the Locator application and searches for all end stations that are attached to the selected switch. All end stations found are displayed in the Locator application's Browse Screen.
- **Webpage** - Opens up a Web session with the selected device. The web session application varies depending on the device. For example, AOS devices will open a WebView session.
- **Resource Manager - Device Inventory** - Launches the Inventory Screen in the Resource Manager Application for the selected switches, which enables you to create and Inventory Report for the selected devices.
- **Resource Manager - Backup Device** - Launches the Backup Wizard in the Resource Manager Application, which enables you to perform a configuration backup of the selected devices.
- **CLI Scripting - Telnet** - Opens up a Telnet session with the selected device in the CLI Scripting application.
- **Notifications - Configure Traps** - Launches the Trap Configuration Wizard in the Notifications application to enable you to configure traps for the selected devices.
- **Notifications - View Traps** - Opens the Notifications Home page to display traps for the selected device.
- **Reboot** - Reboots the selected device(s) You have the option of rebooting from the Working, Certified, or Other Directory and setting a time for the reboot. Click on the Reboot operation link and use the **Reboot From** drop-down to select the directory you want to reboot from. In the **Reboot**

Delay drop-down select when you want to reboot to occur (now, a specific number of minutes from now, or at a specific date and time). Note that when you reboot multiple devices, there is a minimum delay of 30 seconds before the devices reboot (even if you select the Reboot now option). If you select a large number of devices, the delay is equal to roundoff of $(30 + (\text{deviceCount}/4))$, in seconds (e.g., if you select 1,000 devices, the delay is 280 seconds, or 4 minutes). The delay allows time to push the "Reboot" command to all devices.

- **Copy Running/Working to Certified** - Copies the contents of the working/running directory in the primary CMM to the certified directory in the primary CMM. Note that the Copy Working to Certified command also automatically synchronizes the switch's CMMs after the copy operation is completed.
- **Copy Certified to Working/Running** - Copies the contents of the certified directory in the primary CMM to the working/running directory in the primary CMM.
- **Save to Running** - Saves the primary CMM's current running configuration to the current running directory of the switch. OmniVista supports the Multiple Working Directories Feature on certain devices (e.g., OS10K, OS6900). This feature allows the user to create multiple "working" directories on the switch that can be used to save specific switch configurations. When the Save to Running Command is executed, the device(s) save the CMM's current running directory to the current user-defined "working" directory (Running Directory). Note that if you select a group of devices and some do not support multiple working directories, the devices will save the CMM's current running directory to the device's current "working" directory, whether it is a user-defined directory or the Working Directory.

Discovery Displays

By default, a table containing all discovered device is displayed (Inventory List). You can also click on the Chart View icon  to view graphical charts breaking down discovered device by device type, AOS version, and physical location.

Note: The information displayed in the Inventory List is updated based on the frequency settings configured on the Discovery Setting Frequencies Screen. You can perform an immediate poll on a device(s) to update information by selecting the device(s) in the Inventory List and clicking on the **Rediscover** button at the top of the screen.

Inventory List

The Discovery Screen Inventory List displays basic information on all discovered network devices. Click on a device to display detailed information for the device.

Basic Information

- **Friendly Name** - User-configured name for the device.
- **Name** - The name of the device.
- **Address** - The address of the device.
- **DNS Name** - The DNS name of the device.
- **Type** - The type of device chassis (e.g., OS6850-24).
- **Version** - The version number of the device software (e.g., 6.6.5.96.R02). OmniVista may not be able to determine the software version on some third-party devices. In these cases, the field will be blank.
- **Location** - The physical location of the device (e.g., Test Lab).
- **FTP User Name** - The CLI/FTP user name for the device.
- **SNMP Version** - The SNMP version used to communicate with the device.
- **v1/2 Read Community** - The device's SNMP v1/2 "get" community name, if applicable.

- **v1/2 Write Community** - The device's SNMP v1/2 "set" community name, if applicable.
- **v3 User Name** - The device's SNMP v3 user name, if applicable.
- **Last Upgrade Status** - The status of the last firmware upgrade on the device.
 - "Successful" - Successful BMF and Image upgrade performed.
 - "Successful (BMF)" - Successful BMF upgrade performed.
 - "Successful (Image)" - Successful Image upgrade is performed.
 - "Failed (BMF, Image)" - BMF and Image upgrade failed.
 - "Failed (BMF)" - BMF upgrade failed.
 - "Failed (Image)" - Image upgrade failed.
- **Backup Date** - The date that the device's configuration and/or image files were last backed-up to the OmniVista Server.
- **System Up Time** - The amount of time the device has been "Up" since the last reboot.
- **Backup Version** - The firmware version of the configuration and/or image files that were last backed- up to the OmniVista Server.
- **Last Known Up At** - The date and time when the last poll was initiated on the device.
- **Description** - A description of the device, usually the vendor name and model.
- **Status**- The operational status of the device. It displays "Up" if the device is up and responding to polls. It displays "Down" if the device is down and not responding to polls. It displays "Warning" if the switch has sent at least one warning or critical trap and is thus in the warning state.
- **Traps** - The status of trap configuration for the device. "On" means that traps are enabled. "Off" means that traps are disabled. "Not Configurable" means that traps for this device are not configurable from OmniVista. (Note that traps may have been configured for such devices outside of OmniVista.) "Unknown" means that OmniVista does not know the status of trap configuration on this device.
- **Running From** - For AOS devices, this field indicates whether the switch is running from the Certified directory or from the Working directory. This field is blank for all other devices. For AOS devices, the directory structure that stores the switch's image and configuration files in flash memory is divided into two parts:
 - The Certified directory contains files that have been certified by an authorized user as the default configuration files for the switch. When the switch reboots, it will automatically load its configuration files from the certified directory if the switch detects a difference between the certified directory and the working directory.
 - The Working directory contains files that may or may not have been altered from those in the certified directory. The working directory is a holding place for new files to be tested before committing the files to the certified directory. You can save configuration changes to the working directory. You cannot save configuration changes directly to the certified directory.

Note that the files in the certified directory and in the working directory may be different from the running configuration of the switch, which is contained in RAM. The running configuration is the current operating parameters of the switch, which are originally loaded from the certified or working directory but may have been modified through CLI commands, WebView commands, or OmniVista.

Modifications made to the running configuration must be saved to the working directory (or lost). The working directory can then be copied to the certified directory if and when desired.

Note: OmniVista supports the Multiple Working Directories Feature available on OS10K and OS6900 Switches (AOS Release 7.2.1.R01 and later). This feature allows the user to create multiple Working Directories on the switch that can be used to save specific switch configurations. The user can create any name for these "Working" Directories (e.g., "Marketing Switch 05-23-15"). If the switch is running from one of these user-created directories, the directory name is displayed in this field.

- **Changes** - For AOS devices, this field indicates the state of changes made to the switch's configuration. This field is blank for all other devices. This field can display the following

values:

- **Certified** - Changes have been saved to the working directory, and but the working directory has been copied to the certified directory. The working directory and the certified directory are thus identical.
- **Uncertified** - Changes have been saved to the working directory, but the working directory has not been copied to the certified directory. The working directory and the certified directory are thus different.
- **Unsaved** - Changes have been made to the running configuration of the switch that have not been saved to the working directory.
- **Blank** - When this field is blank for an AOS device, the implication is that OmniVista knows of no unsaved configuration changes and assumes that the working and certified directories in flash memory are identical.
- **Discovered** - The date and time when OmniVista successfully pings or polls the switch for the first time. This value remains unchanged until the switch entry is deleted. This field will remain blank if OmniVista does not ping or poll the switch at all.
- **No. of Licenses Used** - The total number of Core (AOS) or Third-Party licenses being used. For example, a stack of 4 switches would require 4 licenses, a VC of 6 would require 6 licenses. If a stack splits, the number of licenses reserved for the device before the split is maintained even though modules have been reduced to less than 5. This way, the license counts are reserved for the stack to recover.
- **License Type** - The type of license used by the device (e.g., AOS, Third Party).
- **Synchronized Status** - Whether the Primary CMM's working directory is identical to the working directory on the other CMM module (if present). Synchronized - The primary CMM's working directory is identical to the working directory on the secondary CMM.
 - **Not Synchronized** - The primary CMM's working directory is not identical to the working directory on the secondary CMM.
 - **Not Applicable** - Only one CMM module is installed.
 - **Unknown** - The synchronization state is unknown..

Detailed Information

Basic Information

- **Name** - The name of the device.
- **Address** - The address of the device.
- **Status**- The operational status of the device. It displays "Up" if the device is up and responding to polls. It displays "Down" if the device is down and not responding to polls. It displays "Warning" if the switch has sent at least one warning or critical trap and is thus in the warning state.
- **DNS Name** - The DNS name of the device.
- **Type** - The type of device chassis (e.g., OS6850-24).
- **Version** - The version number of the device software (e.g., 6.6.5.96.R02). OmniVista may not be able to determine the software version on some third-party devices. In these cases, the field will be blank.
- **Location** - The physical location of the device (e.g. Test Lab).
- **Description** - A description of the device, usually the vendor name and model.

Security Information

- **FTP User Name** - The user name that OmniVista will use to establish CLI/FTP sessions with the discovered devices. The user name specified will be used to auto-login to devices when CLI sessions are established. It will also be used to perform FTP with the device when configuration files are saved and restored.
- **SNMP Version** - The SNMP version that OmniVista will use to communicate with the device. The default version for AOS devices is v2, but v1 and v3 are also supported.
- **Read Community** - The device's "get" community name. The "get" community name enables OmniVista to read information from the device.
- **Write Community** - The device's "set" community name. The "set" community name enables OmniVista to write information to the device.

License Information

- **No. of Licenses Used** - The total number of Core (AOS) or Third-Party licenses being used. For example, a stack of 4 switches would require 4 licenses, a VC of 6 would require 6 licenses. If a stack splits, the number of licenses reserved for the device before the split is maintained even though modules have been reduced to less than 5. This way, the license counts are reserved for the stack to recover.
- **License Type** - The type of license used by the device (e.g., AOS, Third Party).


Status Information

- **Traps** - The status of trap configuration for the device. "On" means that traps are enabled. "Off" means that traps are disabled. "Not Configurable" means that traps for this device are not configurable from OmniVista. (Note that traps may have been configured for such devices outside of OmniVista.) "Unknown" means that OmniVista does not know the status of trap configuration on this device.
- **Running From** - For AOS devices, this field indicates whether the switch is running from the Certified directory or from the Working directory. This field is blank for all other devices.
- **Changes** - For AOS devices, this field indicates the state of changes made to the switch's configuration. This field is blank for all other devices. This field can display the following values:
 - **Certified** - Changes have been saved to the working directory, and but the working directory has been copied to the certified directory. The working directory and the certified directory are thus identical.
 - **Uncertified** - Changes have been saved to the working directory, but the working directory has not been copied to the certified directory. The working directory and the certified directory are thus different.
 - **Unsaved** - Changes have been made to the running configuration of the switch that have not been saved to the working directory.
 - **Blank** - When this field is blank for an AOS device, the implication is that OmniVista knows of no unsaved configuration changes and assumes that the working and certified directories in flash memory are identical.
- **Last Upgrade Status** - The status of the last firmware upgrade on the device.
 - "Successful" - Successful BMF and Image upgrade performed.
 - "Successful (BMF)" - Successful BMF upgrade performed.
 - "Successful (Image)" - Successful Image upgrade is performed.
 - "Failed (BMF, Image)" - BMF and Image upgrade failed.
 - "Failed (BMF)" - BMF upgrade failed. "Failed (Image)" - Image upgrade failed.

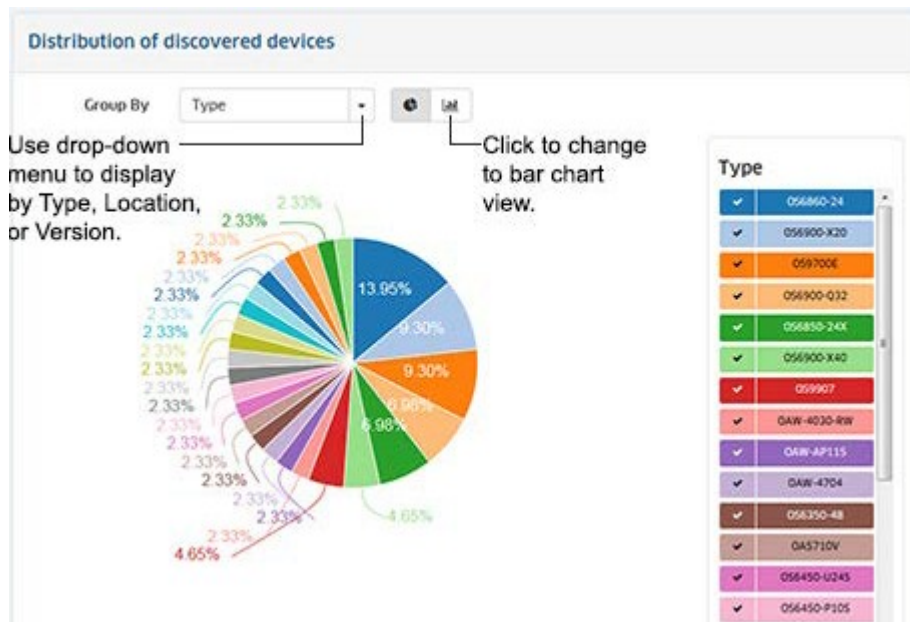
Other Information

- **Backup Date** - The date that the device's configuration and/or image files were last backed-up to the OmniVista Server.
- **Backup Version** - The firmware version of the configuration and/or image files that were last backed-up to the OmniVista Server.
- **System Up Time** - The amount of time the device has been "Up" since the last reboot.
- **Last Known Up At** - The date and time when the last poll was initiated on the device.
- **Discovered** - The date and time when OmniVista successfully pings or polls the switch for the first time. This value remains unchanged until the switch entry is deleted. This field will remain blank if OmniVista does not ping or poll the switch at all.

Graphical Views

For a graphical view of discovered devices grouped by category, click on the Chart View icon  at the top of the screen. By default, the pie chart view is shown, with the inventory information displayed by type. Click on the bar chart option to view the information in bar chart format. Hover the mouse over a section for the number of devices in the category. Change the view using the **Group by** drop-down menu:


- **Type** - Group discovered devices by device type (e.g., OS6860-48, Aruba AP).
- **Location** - Group discovered devices by physical location listed for the device (e.g., NMS Lab, SQA Lab).
- **Version** - Group discovered devices by software version running on the device (e.g., 6.4.3.575.R01, 1.7.1.10).



Discovery Profiles

The Discovery Profiles Screen displays all configured Discovery Profiles and is used to create, edit, and delete profiles. A Discovery Profile is used when discovering network devices. A Discovery Profile contains the parameters that are used by OmniVista when performing a discovery (e.g., SNMP version used to discover devices, CLI/FTP passwords needed to connect to a device).

Creating a Discovery Profile

Click on the Create icon  to bring up the Create Discovery Profile Screen. Complete the fields in each section as described below, then click on the **Create** button.

General

- **Name** - The profile name.
- **CLI/FTP User Name** - The CLI Scripting (Telnet)/FTP user name that OmniVista will use to establish CLI Scripting and FTP sessions with the discovered devices.
- **CLI/FTP Password** - The CLI scripting (Telnet)/FTP user name that OmniVista will use to establish CLI Scripting and FTP sessions with the discovered devices. Note that the user name and password specified will be used to auto-login to devices when CLI Scripting sessions are established. They will also be used to perform FTP with the device when configuration files are saved and restored.
- **Confirm CLI/FTP Password** - Re-enter the CLI/FTP Password.
- **Secondary Password** - Optional Secondary Password used to connect to devices.
- **Confirm Secondary Password** - Re-enter the Secondary Password.

Note: If you do **not** define the CLI/FTP user name and password, and you attempt to save, restore, or upgrade configuration files for AOS devices, you will be individually queried for the FTP login name and password of each individual switch for which configuration files are being saved, restored, or upgraded. In addition, OmniVista will be unable to auto-login to the device when establishing CLI Scripting sessions.

SNMP

- **SNMP Version** - The SNMP version used to discover devices (v1, v2, v3). (Default = v2)
- **Timeout** - The time period, in milliseconds, that OmniVista will wait for a switch to respond to a connection request before assuming that the request has timed-out.
- **Read Community (v1 and v2 only)** - The Read Community Name, which is used to read information from a device.
- **Write Community (v1 and v2 only)** - The Write Community Name, which is used to write information to a device.
- **Retry Count** - The number of times OmniVista will attempt to attempt to connect to a switch.
- **User Name (v3 only)** - The SNMP version 3 user name.
- **Auth and Priv Protocol (v3 only)** - Select the authentication protocol OmniVista will use for SNMP communications with the discovered switches (None, MD5, or SHA).
- **Auth Password (v3 only)** - The password that OmniVista will use for the MD5 or SHA authentication protocol (if applicable).
- **Confirm Auth Password (v3 only)** - Confirm the authentication password entered above.
- **Priv Password (v3 only)** - The password that will be used as the secret key (if applicable).
- **Confirm Priv Password (v3 only)** - Confirm the privilege password entered above.
- **Context Name (v3 only)** - The unique context name for this context. (An SNMP context is a collection of management information accessible by an SNMP entity, in this case OmniVista.)
- **Context ID (v3 only)** - The unique context ID for this context. Each context must be identified by a unique context name and a unique context ID.


Note: If a device's Read and Write Community Names are "public" (Default), you can leave these fields blank (OmniVista uses the default name, "public" when the field is blank.) Read and Write Community Names are not configurable from OmniVista; they can only be configured by logging onto a device. Also note that when you use SNMP v3, Read and Write Community

Names are ignored.


Advanced Settings

- **Trap Station Name** - The device user name that will be used when an AOS device is configured to send traps to OmniVista. AOS devices require that a valid device user name be specified with the trap station configuration entry. If this field is left blank, the following switch user names will be used by default for trap station configuration entries:
 - If OmniVista is configured to use SNMP version 3 with this device, the SNMP version 3 user name entered for the device will be used as the switch user name in the trap station configuration entry.
 - If OmniVista is configured to use SNMP version 1 or SNMP version 2 with this device, the read community string for the device will be used as the switch user name in the trap station configuration entry.
- **Discover Link** - Specifies how OmniVista will discover the physical links associated with the discovered devices. Links to other switches are displayed graphically on OmniVista's Topology maps.
 - **Normally** - This setting is used for devices that support adjacency protocols, such as AOS devices. Adjacency protocols (such as XMAP and AMAP) enable OmniVista to discover the physical links associated with specific devices.
 - **As OEM Device** - This setting enables you to use the new "end station search" functionality from the Locator application to automatically discover links for devices that do not support adjacency protocols. If this setting is used and the device does not support an adjacency protocol that enables OmniVista to discover physical links, the end station search algorithms used by the Locator application are invoked at each polling cycle to discover the device's links. All links discovered are displayed on Topology maps automatically.
- **Shell Preference** - Specifies the default command line interface to be used for discovered devices. OmniVista's CLI Scripting application supports both the Telnet and SSH command line interfaces. SSH (Secure Shell) is a Telnet-like utility that provides encryption and is far more secure than Telnet. When the SSH setting is used, SSH will be used as the default command line interface for the device. If the Telnet setting is used, Telnet will be used as the default command line interface for the device. Ensure that devices are capable of SSH before you use the SSH setting. OmniVista does not verify devices' SSH capabilities. All AOS devices are SSH-capable. (Default = SSH)
- **Use Get Bulk** - Enables (Yes)/Disables (No) the "Get Bulk" operation. When enabled, the "Get Bulk" operation is used for retrieving large amounts of data, particularly from large tables. The Get Bulk operation performs continuous "Get Next" operations, each time requesting the number of table rows specified by the value in the **Max Repetitions** field (described below). For example, if the value in the **Max Repetitions** field is ten, each Get Next operation will request 10 rows of table data. Note that the number of rows of data actually returned by the switch will be determined by the amount of memory the switch has available at that time.
- **Max Repetitions** - The number of rows of table data that the "Get Bulk" operation will request in each "Get Next" operation.

Editing a Discovery Profile

Select a profile from the Existing Profiles Table and click on the Edit icon . Update any fields as described above and click on the **Update** button. Note that you cannot edit a profile name.

Deleting a Discovery Profile

Select a profile(s) from the Existing Profiles Table, click on the Delete icon , then click **OK** at the confirmation prompt.

Profile Information

Basic Discovery Profile information is displayed in the Existing Profiles Table. Click on a profile to display detailed information.

Basic Information

- **Name** - The profile name.
- **SNMP Version** - The SNMP version used to discover devices (v1, v2, v3). (Default = v2)
- **v1/2 Read Community** - The SNMP v1/v2 Read Community Name, which is used to read information from a device, if applicable.
- **v1/2 Write Community** - The SNMP v1/v2 Write Community Name, which is used to write information to a device, if applicable.
- **v3 User Name** - The SNMP v3 user name, if applicable.

Detailed Information

- **Name** - The profile name.
- **CLI/FTP User Name** - The CLI Scripting (Telnet)/FTP user name that OmniVista will use to establish CLI Scripting and FTP sessions with the discovered devices.
- **SNMP Version** - The SNMP version used to discover devices (v1, v2, v3). (Default = v2)
- **Timeout** - The time period, in milliseconds, that OmniVista will wait for a switch to respond to a connection request before assuming that the request has timed-out.
- **Read Community (v1 and v2 only)** - The Read Community Name, which is used to read information from a device.
- **Write Community (v1 and v2 only)** - The Write Community Name, which is used to write information to a device.
- **Retry Count** - The number of times OmniVista will attempt to attempt to connect to a switch.
- **User Name (v3 only)** - The SNMP version 3 user name.
- **Auth and Priv Protocol (v3 only)** - The authentication protocol OmniVista will use for SNMP communications with the discovered switches (None, MD5, or SHA).
- **Auth Password (v3 only)** - The password that OmniVista will use for the MD5 or SHA authentication protocol (if applicable).
- **Confirm Auth Password (v3 only)** - Confirm the authentication password entered above.
- **Priv Password (v3 only)** - The password that will be used as the secret key (if applicable).
- **Confirm Priv Password (v3 only)** - Confirm the privilege password entered above.
- **Context Name (v3 only)** - The unique context name for this context. (An SNMP context is a collection of management information accessible by an SNMP entity, in this case OmniVista.)
- **Context ID (v3 only)** - The unique context ID for this context. Each context must be identified by a unique context name and a unique context ID.
- **Trap Station User Name** - The device user name that will be used when an AOS device is configured to send traps to OmniVista. AOS devices require that a valid device user name be specified with the trap station configuration entry. If this field is left blank, the following switch user names will be used by default for trap station configuration entries:
 - If OmniVista is configured to use SNMP version 3 with this device, the SNMP version 3 user name entered for the device will be used as the switch user name in the trap station configuration entry.
 - If OmniVista is configured to use SNMP version 1 or SNMP version 2 with this device, the read community string for the device will be used as the switch user name in the trap station

configuration entry.

- **Discover Link** - Specifies how OmniVista will discover the physical links associated with the discovered devices. Links to other switches are displayed graphically on OmniVista's Topology maps:
 - **Normally** - This setting is used for devices that support adjacency protocols, such as AOS devices. Adjacency protocols (such as XMAP and AMAP) enable OmniVista to discover the physical links associated with specific devices.
 - **As OEM Device** - This setting setting enables you to use the new "end station search" functionality from the Locator application to automatically discover links for devices that do not support adjacency protocols. If this setting is used and the device does not support an adjacency protocol that enables OmniVista to discover physical links, the end station search algorithms used by the Locator application are invoked at each polling cycle to discover the device's links. All links discovered are displayed on Topology maps automatically.
- **Shell Preference** - Specifies the default command line interface to be used for discovered devices. OmniVista's CLI Scripting application supports both the Telnet and SSH command line interfaces. SSH (Secure Shell) is a Telnet-like utility that provides encryption and is far more secure than Telnet. When the SSH setting is used, SSH will be used as the default command line interface for the device. If the Telnet setting is used, Telnet will be used as the default command line interface for the device. (Default = SSH)
- **Use Get Bulk** - Enables (Yes)/Disables (No) the "Get Bulk" operation. When enabled, the "Get Bulk" operation is used for retrieving large amounts of data, particularly from large tables. The Get Bulk operation performs continuous "Get Next" operations, each time requesting the number of table rows specified by the value in the **Max Repetitions** field (described below). For example, if the value in the **Max Repetitions** field is ten, each Get Next operation will request 10 rows of table data. Note that the number of rows of data actually returned by the switch will be determined by the amount of memory the switch has available at that time.
- **Max Repetitions** - The number of rows of table data that the "Get Bulk" operation will request in each "Get Next" operation.

Third-Party Devices Support

The Discovery Third-Party Devices Support Screen is used to enable discovery and support of third-party devices. The Third-Party Devices Support Screen enables you to add support for third-party devices, edit third-party device support, delete support for unwanted third-party devices. The Mibset List displays all configured third-party device support information.

Note: Support for Cisco and Extreme devices must be added manually as described below.

Adding Third-Party Device Support

To add support for a third-party device, click on the Add icon **+** and complete the fields as described below. When you have completed the fields, click on the **Create** button. The entry will appear in the MibsetList.

- **IOD** - The device's Object ID. Enter only the portion of the OID relative to the ".1.3.6.1.4.1." (".iso.org.dod.internet.private.enterprises") branch. For example, enter only '9' for Cisco devices rather than '.1.3.6.1.4.1.9', or '1916' for Extreme devices, rather than '.1.3.6.1.4.1.1916'. Using this vendor value (e.g., 9, 1916) will enable OmniVista to recognize all devices from the vendor. Note that you can also enter specific vendor device values (e.g., '1916.800.1.1.2.1.5.1') for each vendor device if you want each device to have a different name while using the same mibset.
- **Display Name** - The name to be used for the device.
- **MIB Directory Name** - The directory name of the device's MIB. If you want to use MIB-2 level

support for third-party devices, enter **mib-2**. This generic MIB-2 directory already exists in OmniVista. If you are not using standard MIB-2, enter a new directory name for the MIB. Note that the directory does not have to actually exist; it will be created automatically when you import the MIB.

- **Enabled** - Select **On** or **Off** to enable (On) or disable (Off) discovery for the device.
- **Icon** - The generic third-party icon appears in the Icon field. If you have an icon you would like to display for the device, click on the **Choose Image** button and locate the image. The image will appear in the Icon field.

Traps for Third-Party Devices


By default, OmniVista supports generic MIB-2 traps for third-party devices. If you import a new, custom MIB for a third-party device, OmniVista will automatically scan the MIB for new traps and integrate any traps it finds. Note that MIBs do not include synopses for traps. OmniVista will create a synopsis "on the fly" for any new trap it integrates. You can go to the Trap Definition Screen in the Notifications application and edit the synopses or severity levels that OmniVista assigns to new traps.

Third-Party Device Support After Discovery

Once third-party devices have been discovered, OmniVista supports the following functionality for the devices:

- **Web Browser** - OmniVista enables you to launch web-based element managers for third-party devices using the "Webpage" operation in the Topology application.
- **Telnet or SSH (as applicable)** - OmniVista enables you to initiate Telnet or SSH sessions to third-party devices using the Terminal Screen in the CLI Scripting application.
- **Custom MIBs** - OmniVista allows you to import custom MIBs for third-party devices (as described above).
- **Custom Icons** - OmniVista enables you to import a custom icon that will be used to represent a specific third-party device.
- **Traps** - By default, OmniVista supports generic MIB-2 traps for third-party devices. In addition, whenever you import a new, custom MIB for a third-party device, OmniVista will scan the MIB for new traps and automatically integrate any traps it finds.
- **Locator** - OmniVista's Locator application supports third-party devices.

Editing Third-Party Device Support

Select an entry in the Mibset List and click on the Edit icon . Edit the fields as described above, then click on the **Update** button. Note that you cannot edit the OID or Display Name.

Deleting Third-Party Device Support

Select an entry in the Mibset List and click on the Delete icon . Click **OK** at the confirmation prompt.

Mibset List

The Mibset List displays information about configured third-party device support.

- **OID** - The device's Object ID.
- **Display Name** - The name that is used for the device.
- **MIB Directory Name** - The directory that contains the device's MIB.
- **Enabled** - "True" (enabled) indicates that the device is included in the discovery process. "False" (disabled) indicates that the device is not included in the discovery process.

Import MIBs

The Discovery Import MIBs Screen is used to import new or updated MIB files into OmniVista. All MIB files are imported to the OmniVista Server. Before you import MIBs, it is important to understand that the purpose of this function is to import MIB files that reside somewhere on your local file system into OmniVista. A mibs.txt ASCII file lists the order in which the MIBs will be compiled. Also:

- All MIB files that you import must have a file extension of **.mib**.
- If you create a new MIB directory for a new device, note that you must import a complete set of MIBs into that directory. This means that if any proprietary MIBs you are using have imports of standard MIBs, the standard MIBs must be included and imported into that directory as well.
- For the MIBs to compile correctly, you are strongly advised to order them so that all the referenced MIB files are compiled before the files that reference them. MIB compilers follow import references from one MIB to another on the fly, and do not strictly require that the MIBs be compiled in any particular order.

For this to work successfully, however, the MIB filenames must match the import statements exactly, and unfortunately this is almost never the case. To avoid these problems, as stated above, order the MIB files so that all the referenced MIB files are compiled before the files that reference them. You can specify the order in which the MIB files will be compiled by selecting files and using the **Up** and **Down** arrows in the Import Files to Mibset Screen, as described in the procedure below. MIB files will be compiled in the order that the files are listed in the Import Files to Mibset Screen.

- It is not advisable to add new MIB files to a MIB directory supplied by default with OmniVista. It is preferable to create a separate new directory for each new third-party device you want to support. This will ensure proper operation of the OmniVista MIB Browser. If you add a new MIB file to an existing MIB directory, you will need to re-import the existing MIB files in order for them all to display in the OmniVista MIB Browser.
- Once you have completed the MIB importation process, OmniVista does not immediately parse the MIBs. When you discover a device with an OID that is specified for the MIB directory into which you imported the new MIBs, OmniVista will poll the device for standard MIB-II objects. If the standard MIB-II MIBs are not included in the directory, error messages will be written to file server.txt (which can be viewed from the Audit application). Any proprietary MIBs that you imported into the directory will not be parsed until you load the MIB Browser for a device with an OID that is specified for that directory.
However, if you close the OmniVista client and completely stop the OmniVista server after completing the MIB importation process, then start the server, the MIBs will be parsed when the server starts.

Importing MIBs

Follow the steps below to import MIB files into OmniVista.

1. Select the **Mibset to be updated** from the drop-down box at the top of the screen (e.g., Cisco). If you entered a new directory name in the Third-Party Device Support Screen, the name will be displayed in the drop-down menu.
2. Click on the **Import** button, then click on the **Upload Files** button.
3. Browse to the folder containing the MIBs you want to import, select all of the files and click **Open**. The files will appear in the imported into the Import Files to Mibset Screen.

Note: If you are using Chrome, you will have the option of selecting an **Upload Folder** button in Step 2. Select the folder containing the MIB Files to import all of the files in the folder. This option is not supported in Firefox or Internet Explorer.

4. ¶ The MIB files will be loaded into OmniVista in the correct order. However, you can adjust this order by selecting individual files and clicking the **Up** and **Down** arrows in the upper-right corner of the screen.

5. Click the **Apply** button. The MIB files are imported to the OmniVista Server.

Hardware Inventory

The Discovery Hardware Inventory Screen is used to view hardware inventory information (e.g., CMM, Chassis, Power Supplies) for discovered devices. To view information for a device, select an option from the drop down menu (Use Switch Picker/Use Topology), click on the **Select Device** button and select a device.

Asset Information

- **Friendly Name** - A user-definable name for the device. If no name was configured, the IP address of the device is displayed.
- **Module Type** - The physical type of module or submodule in this physical location (e.g., Chassis, NI). Note that the value for this field displays as "Unknown" for a brief period while a newly-installed module or submodule is identified.
- **Module Name** - The manufacturer's name for the module (e.g. OS6850--C48, OS6850-BPS-PS).
- **Description** - The user-definable description of this particular module or submodule. The module description can be defined through SNMP.
- **Serial Number** - The serial number of the module or sub-module.
- **Part Number** - The part number of the module or sub-module.
- **MAC Address** - The base MAC address for the module or submodule. If not applicable, the field will be blank.
- **OS Version** - The OS version number running on the module. If not applicable, the field will be blank.
- **Uboot Version** - The U-Boot version running on the module. If not applicable, the field will be blank.
- **HW Revision** - The hardware revision number for the module. If not applicable, the field will be blank.
- **Firmware Version** - The version/revision level of the module or submodule firmware. If not applicable (e.g., Power Supply), the field will be blank.
- **Manufacturer Name** - The manufacturer of the module.
- **License** - Additional licenses (other than the Core License) active on the module (e.g., Advanced), if applicable.
- **Slot** - The slot in the chassis where the module resides. If not applicable, the field will be blank.

Note: Information is displayed for the following devices: OS6250, OS6350, OS6400, OS6450, OS6850/OS6850E, OS6855, OS9000, OS6860/OS6860E, OS6865, OS6900, OS9700/OS9700E, OS9800/OS9800E, OS9900, OS10K, OAW-4xxx (running Alcatel OS only), OAW-AP (running Alcatel OS only).

Ports

The Discovery Ports Screen is used to display information about ports on network devices, and is also used to enable/disable device ports. Click on the **Select Devices** button and use the Switch Picker or Topology option to select the devices you want to view. You can select up to 50 devices at a time. You can also enable/disable ports by selecting the port(s) and click on the **Enable** or **Disable** button at the top of the screen.

Port Information

- **Friendly Name** - Displays the device label as configured in the Preferences application (e.g., device IP address, System name, DNS name).
- **IP Address** - The IP address of the device on which the port resides.
- **Name** -The user-defined name for the device.
- **Slot/Port** - The slot/port number on the device.
- **Port Alias** - The user-defined alias for the port.
- **Port Description** - A detailed description of the interface (e.g., Alcatel-Lucent OS6900 QNI 1/1A)
- **Admin Status** - The administrative status of the port (Up/Down). When the admin status of a port is "Up", the port can receive and transmit data as long as a cable is connected and no physical or operational problems exist. When the admin status of a port is "Down", the port will not transmit or receive data even if a cable is connected and the physical connection is operational. Note that physical or operational problems may cause a port to be nonfunctional even when its admin status is enabled.
- **Operational Status** - The operational status of the port (Up/Down/Unknown). If the operational status of a port is "Up", the admin status of the port is "Up" and a cable is connected to the port and transmitting data.

Note: If an interface's Admin Status is "Down", its Operational Status will also be "Down". When the Admin Status is changed to up, the interface's Operational Status will change to "Up" if the interface is ready to transmit and receive packets; or the Operational Status will change to "Dormant" if the interface is waiting for external actions; or the Operational Status will remain "Down" if there is a fault that prevents it from going "Up".

- **Operational State Changed At** - The date and time the operational status of the port last changed (from "Up" to "Down" or "Down" to "Up").
- **Configured Speed** - The configured interface line speed, in Mbps. If the port is set to "Auto", the switch automatically sets the line speed to match the attached device.
- **Negotiated Speed** - The actual speed of the port.
- **Default VLAN** - The default VLAN to which the port is assigned.
- **Port Type** - The port type (e.g., Ethernet-CSMA/CD).
- **Specific Type** - Detailed information about the port type (e.g., LAG, Stack).
- **Port Properties** - The port configuration properties (e.g., UNP, Mobile, LLDP)
- **UNP Status** - Indicates whether or not UNP is Enabled/Disabled on the port.
- **Port Split Status** - The split port status of the switch (Auto, 40GB, or 4 x 10GB). The Split Port Feature is only supported on OS6900-Q32/X72 Switches running AOS 7.3.4.R01 and later. If the port is on a non-supported device, the field will display "False".
- **VFL Port** - Indicates whether or not the port is a VFL port (True/False)
- **PoE Status** - The administrative status of PoE on the port (True=enabled, False=disabled). If the device/port does not support PoE, the field will display "False".
- **PoE Wattage** - The amount of PoE power being used by the device (if applicable), in milliwatts.


Note: The Discovery Ports feature is supported on the following devices: OS6250, OS6350, OS6400, OS6450, OS6850/OS6850E, OS6855, OS9000, OS6860/OS6860E, OS6865, OS6900, OS9700/OS9700E, OS9800/OS9800E, OS9900, OS10K.

Link

The Discovery Link Screen displays all links that were learned during the discovery process, or created manually in OmniVista. It is also used to manually create, edit, and delete manual links. Unlike automatically- discovered Links, which disappear from the Topology map view when they become


unreachable, manual links will be persistent and display in RED when the link goes down. This enables users to manually configure critical links, such as the network core links (which are seldom changed), providing better monitoring capability for critical links.

Creating a Link


Click on the Create icon  and complete the fields as described below. When you are finished, click on the **Create** button.

- **IP Address 1** - The IP address of one device in the link. Select an option from the drop down menu (Use Switch Picker/Use Topology), click on the **Select Device** button and select a device.
- **Slot/Port 1** - Select a slot/port for IP address 1 from the drop-down menu.
- **LAG 1** - If this is a link aggregation link, set the LAG 1 field to the Link Aggregation Number assigned by the device above when the link aggregation group was created.
- **IP Address 2** - The IP address of the second device in the link. Select an option from the drop down menu (Use Switch Picker/Use Topology), click on the **Select Device** button and select a device.
- **Slot/Port 2** - Select a slot/port for IP address 2 from the drop-down menu.
- **LAG 2** - If this is a link aggregation link, set the LAG 2 field to the Link Aggregation Number assigned by the device above when the link aggregation group was created.
- **Media Type** - Select the media type for the link from the drop-down menu.
- **Status** - Select the administrative status for the link from the drop-down menu (Up/Down). Note that you can edit the link later if you want to change the status.

Editing a Link

Select a link in the Existing Links Table and click on the Edit icon . Edit the fields as described above and click on the **Apply** button.

Deleting a Link

Select a link in the Existing Links Table and click on the Delete  icon . Click **OK** at the Confirmation Prompt.

Existing Links Table

- **Origin** - The origin of the link (e.g., AMAP, LLDP, Manual).
- **IP Address 1** - The IP address of one switch in the link.
- **Slot/Port 1** - The slot and port that connect the link on IP address 1.
- **LAG 1** - If this is a link aggregation link, this field displays the Link Aggregation reference number assigned by the first switch when the link aggregation group was created.
- **IP Address 2** - The slot and port that connect the link on IP address 1.
- **Slot/Port 2** - The slot and port that connect the link on the second switch, specified above.
- **LAG 2** - If this is a link aggregation link, this field displays the Link Aggregation reference number assigned by the second switch when the link aggregation group was created.
- **Ring ID** - The Ethernet Ring Protection (ERP) ID, if applicable.
- **Media Type** - The media type of the link (e.g., Ethernet).
- **Status** - The status of the link (e.g., Up/Down/Unknown).

Setting Frequencies

Once the first discovery is complete, OmniVista performs automatic periodic discoveries to keep its information about the network updated. The Discovery Setting Frequencies Screen is used to configure the frequency of automatic periodic discoveries.

Configuring Automatic Polling

You can configure the frequency of automatic polling. Enter a value (Days, Hours, Minutes) in the applicable Automatic Discovery Type and click on the **Apply** button.

Automatic Discovery Types

OmniVista performs the following automatic discoveries:

- Full Discovery
- Occasional Updates
- Regular Updates
- Frequent Updates

Full Discovery

By default, OmniVista makes a Full Discovery once every eight (8) hours. Full Discoveries include:

- Down Switch Polling
- Frequent Update Polling
- Regular Update Polling
- Auto-discovery of network devices as specified in the Discovery application.

Occasional Updates

By default, OmniVista makes Occasional Updates once every four (4) hours. Occasional Updates include:
Down Switch Polling

- Frequent Update Polling
- Regular Update Polling

Regular Updates

By default, OmniVista makes Regular Updates once every hour. Regular Updates include:

- Down Switch Polling
- Frequent Update polling as described above. Additional polling for:
 - Detailed chassis, module, and port information
 - VLAN information
 - Link Aggregation
 - Ethernet link discovery (i.e., polling AMAP tables) Locator:
- MAC address column from the ARP Table
- Bridge Forwarding Table

Frequent Updates

By default, OmniVista makes Frequent Updates every five (5) minutes. Frequent Updates include: Down Switch Polling

- Polling the standard MIB-II scalar variables sysName and sysDescr
-

- For AOS devices, polling for:
 - The running directory (certified or working), the certification status, and the administrative status of all CMMs.
 - The configuration change status; i.e., has the configuration changed since the last save of memory.

IP Failover

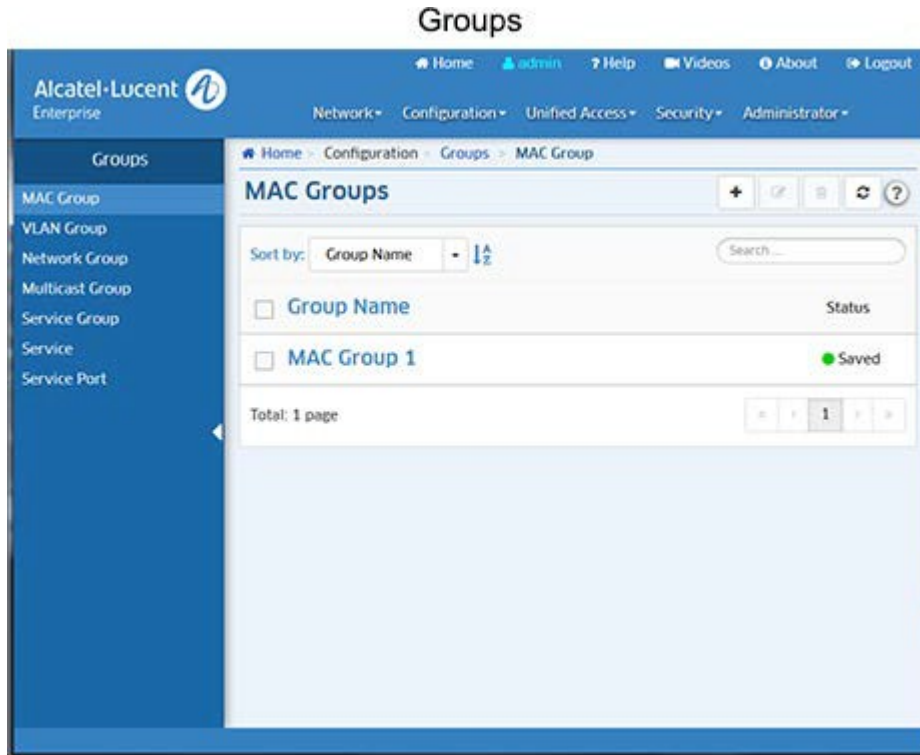
The Discovery IP Failover Screen is used to specify whether or not OmniVista will use a device's alternate IP address for SNMP traffic if the primary IP address fails. If IP Failover is enabled and a device fails to respond to SNMP requests, the OmniVista Server tries to reach the switch using the alternate IP address. If the attempt is successful, all subsequent management traffic is diverted to this new address. Use the **IP Failover** slider to enable (On) or disable (Off) the feature and click on the **Apply** button.

Switch Monitoring

The Discovery Switch Monitoring Screen is used to configure discovery polling of "down" devices. OmniVista polls down switches once per minute to check if the switches have come back up. Select the **Always** radio button if you want this monitoring to occur all the time. Select the **Only if Polling Enabled** radio button if you want this monitoring to occur only when normal OmniVista polling is enabled. After making a selection, click the **Apply** button. The change takes effect immediately.

12.0 Groups

The Groups Application enables you to create groups, which can be used in various PolicyView conditions. Groups are stored on an LDAP (Lightweight Directory Access Protocol) repository that is automatically installed with OmniVista and resides on the same device as the OmniVista server. When the switches in the network are notified to re-cache their policy information, the firmware loads the groups referred by these policies. The Groups application enables you to create five types of groups: MAC Groups, VLAN Groups, Network Groups, Multicast Groups, and Service Groups. You can also configure Services and Service Ports to be used in Service Groups.



MAC Groups





The Groups MAC Groups Screen displays all configured MAC Groups. The screen is used to create, edit, and delete MAC Groups, which can be used in creating various policy conditions, such as source MAC group condition and destination MAC group condition.

Creating a MAC Group


Click on the Create icon **+**. Enter a **Name** for the MAC Group. Enter a **MAC Address** and click on the Add icon **+**. Repeat to add additional addresses. When you are done, click on the **Create** button. The MAC Group will appear in MAC Groups List. Note that you must enter at least one MAC Address.

Editing a MAC Group

Click on the MAC Group that you want to edit to view the MAC Addresses in the MAC Group. Note that you cannot edit a MAC Group name. To edit a MAC Group name you must delete the MAC Group and create a new one.

- To **add** a MAC Address to the Group, enter the **MAC Address**, then click on the Add icon . Repeat to add additional addresses. When you are done, click on **Update** button.
- To **edit** a MAC Address, click on the Edit icon , edit the address, then click on the Save icon . Repeat to edit additional addresses. When you are done, click on the **Update** button.
- To **delete** a MAC Address, click on the Delete icon  next to the MAC Address you want to delete. Repeat to delete additional addresses. When you are done, click on the **Update** button.

Deleting a MAC Group



To delete a MAC Group(s), select the checkbox next to the group(s) in the list, click on the Delete icon , then click **OK** at the confirmation prompt.

Note: MAC Groups that are in use by policy conditions cannot be deleted. To delete these MAC groups, remove them from the policy conditions.

VLAN Groups





The Groups VLAN Groups Screen displays all configured VLAN Groups. The screen is used to create, edit, and delete VLAN Groups.

Creating a VLAN Group


Click on the Create icon . Enter a **Name** for the VLAN Group. Enter a **VLAN Range** and click on the Add icon . Repeat to add additional VLAN Ranges. When you are finished, click on the **Create** button. The VLAN Group will appear in VLAN Groups List. Note that you must enter at least one (1) VLAN range.

Editing a VLAN Group

Click on the VLAN Group that you want to edit to view the VLAN ranges in the VLAN Group. Note that you cannot edit a VLAN Group name. To edit a VLAN Group name you must delete the VLAN Group and create a new one.

- To **add** a VLAN Range to the Group, enter a **VLAN Range** and click on the Add icon . Repeat to add additional ranges. When you are done, click on the **Update** button.
- To **edit** a VLAN Range, click on the Edit icon , edit the range, then click on the Save icon . Repeat to edit additional ranges. When you are done, click on the **Update** button.
- To **delete** a VLAN Range, click on the Delete icon  next to the VLAN Range you want to delete. Repeat to delete additional ranges. When you are done, click on the **Update** button.

Deleting a VLAN Group



To delete a VLAN Group(s), select the checkbox next to the group(s) in the list, click on the Delete icon , then click **OK** at the confirmation prompt.

Note: VLAN Groups that are in use by policy conditions cannot be deleted. To delete these VLAN groups, remove them from the policy conditions.

Network Groups





The Groups Network Groups Screen displays all configured Network Groups. The screen is used to create, edit, and delete Network Groups.

Creating a Network Group


Click on the Create icon . Enter a **Name** for the Network Group. Enter a **Subnet IP/Subnet Mask** and click on the Add icon . Repeat to add additional subnets. When you are finished, click on the **Create** button. The Network Group will appear in Network Groups List. Note that you must enter at least one Subnet IP/Subnet Mask.

Editing a Network Group

Click on the Network Group that you want to edit to view the Subnets in the Network Group. Note that you cannot edit a Network Group name. To edit a Network Group name you must delete the Network Group and create a new one.

- To **add** a Subnet Address to the Group, enter a **Subnet IP/Subnet Mask** and click on the Add icon . Repeat to add additional subnets. When you are finished, click on the **Update** button.
- To **edit** a Subnet, click on the Edit icon , edit the address, then click on the Save icon . Repeat to edit additional Subnets. When you are done, click on the **Update** button.
- To **delete** a Subnet, click on the Delete icon  next to the Subnet you want to delete. Repeat to delete Subnets. When you are done, click on the **Update** button.

Deleting a Network Group




To delete a Network Group(s), select the checkbox next to the group(s) in the list, click on the Delete icon , then click **OK** at the confirmation prompt.

Note: Network Groups that are in use by policy conditions cannot be deleted. To delete these Network groups, remove them from the policy conditions.

Multicast Groups





The Groups Multicast Groups Screen displays all configured Multicast Groups. The screen is used to create, edit, and delete Multicast Groups.

Creating a Multicast Group


Click on the Create icon . Enter a **Name** for the Multicast Group. Enter a **Subnet IP/Subnet Mask** and click on the Add icon . To add additional subnets, click on the Add icon  and enter the subnets. When you are finished, click on the **Create** button. The Multicast Group will appear in Multicast Groups List. Note that you must enter at least one Subnet IP/Subnet Mask. .

Editing a Multicast Group

Click on the Multicast Group that you want to edit to view the Subnets in the Multicast Group.

- To **add** a Subnet Address to the Group, enter the **Subnet IP/Subnet Mask**, then click on the Add icon . Repeat to add additional subnets. When you are done, click on the **Update** button.
- To **edit** a Subnet, click on the Edit icon , edit the address, then click on the Save icon . Repeat to edit additional Subnets. When you are done, click on the **Update** button.
- To **delete** a Subnet, click on the Delete icon  next to the Subnet you want to delete. Repeat to delete Subnets. When you are done, click on the **Update** button.

Deleting a Multicast Group


To delete a Multicast Group(s), select the checkbox next to the group(s) in the list, click on the Delete icon , then click **OK** at the confirmation prompt.

Note: Multicast Groups that are in use by policy conditions cannot be deleted. To delete these Multicast groups, remove them from the policy conditions.

Service Groups

The Groups Service Groups Screen displays all configured Service Groups. The screen is used to create, edit, and delete Service Groups.


Creating a Service Group

Click on the Create icon **+**. Enter a **Group Name** for the Service Group. Select a Service(s) and click on the **Create** button. If you want to create a new Service, click on the Add Icon  to go to the Services Screen and create the Service. When you click on the **Create** button on the Services Screen you will be returned to the Create Service Group Screen to finish creating the Service Group. Note that you must enter at least one service. Also, you cannot use Source and Destination Services in group.

Editing a Service Group

Click on the Service Group that you want to edit, then click on the Edit Icon . Add or remove Services from the group as described above then click on the **Update** button. You cannot edit a Service Group name. To edit a Service Group name you must delete the Service Group and create a new one.

Deleting a Service Group

To delete a Service Group(s), select the checkbox next to the group(s) in the list, click on the Delete icon , then click **OK** at the confirmation prompt.


Note: Service Groups that are in use by policy conditions cannot be deleted. To delete these Service Groups, remove them from the policy conditions.

Services


The Groups Services Screen displays all configured Services, which are used to create Service Groups. The screen is used to create, edit, and delete Services.

Creating a Service


Click on the Create icon **+**. Complete the fields as described below, then click on the **Create** button.

- **Service Name** - User-configured name for the Service.
- **Protocol** - Select a protocol for the Service. By default, the TCP radio button is selected and TCP ports are displayed. Click on the UDP radio button to display UDP ports.
- **Source Port** - Select a source port from the Source Port drop-down list. The drop-down box includes a list of well-known TCP or UDP ports. Select a port(s) from the drop down menu (you can also select "Check All" to select all ports. Click "Uncheck All" to deselect the ports and start over). If you want to create a new port, click on the Add Icon  to go to the Service Port Screen and create a new port.
- When you click on the **Create** button on the Service Port Screen you will be returned to the


Create Service Screen to finish creating the Service. Note that you can specify a Source Port, a Destination Port, or both.

- **Destination Port** - Select a destination port from the Destination Port drop-down list. The drop-down box includes a list of well-known TCP or UDP ports. Select a port(s) from the drop down menu (you can also select "Check All" to select all ports. Click "Uncheck All" to deselect the ports and start over). If you want to create a new port, click on the Add Icon  to go to the Service Port Screen and create a new port. When you click on the **Create** button on the Service Port Screen you will be returned to the Create Service Screen to finish creating the Service. Note that you can specify a Source Port, a Destination Port, or both.

Editing a Service

Click on the Service that you want to edit, then click on the Edit Icon . Edit the field(s) as described above then click on the **Update** button. You cannot edit a Service Name. To edit a Service Name, you must delete the Service and create a new one.

Deleting a Service


To delete a Service(s), select the checkbox next to the Service(s) in the list, click on the Delete icon , then click **OK** at the confirmation prompt.

Note: Services that are in use by policy conditions cannot be deleted. To delete these Services, remove them from the policy conditions.

Service Port


The Groups Service Port Screen displays all configured Service Ports, which are used to create Services. By default, the TCP radio button is selected and TCP Services are displayed. Click on the UDP radio button to display UDP Services. The screen is used to create, edit, and delete Service Ports.

Creating a Service Port


Click on the Create icon . Complete the fields as described below, then click on the **Create** button.

- **Name** - User-configured name for the Service Port.
- **Port Number** - Enter a Service Port number.

Editing a Service Port

Click on the Service Port that you want to edit, then click on the Edit Icon . Edit the field(s) as described above then click on the **Update** button. You cannot edit a Service Port name. To edit a Service Port name you must delete the Service Port and create a new one.

Deleting a Service Port

To delete a Service Port(s), select the checkbox next to the port(s) in the list, click on the Delete icon , then click **OK** at the confirmation prompt.

Note: Service Ports that are in use by policy conditions cannot be deleted. To delete these Service Ports, remove them from the policy conditions.

13.0 PIM

Protocol-Independent Multicast (PIM) is an IP multicast routing protocol that uses routing information provided by unicast routing protocols. Creation of Multicast VXLAN Services requires PIM configuration on devices in the VXLAN. The following screens are used to configure and view PIM configuration on the network:

- **PIM Global Configuration** - Used to configure PIM configuration profiles, which can be applied to switches on the network.
- **PIM Interface** - Used to configure a PIM interface(s) on a switch(es).
- **PIM Candidate** - Used to configure Candidate Bootstrap Routers (C-BSRs) and Bootstrap Routers (BSRs).
- **PIM Device View** - Used to view PIM configurations on network switches.



PIM Global Configuration

The PIM Global Configuration Screen displays all configured PIM Global Profiles and is used to create, edit, assign, and delete PIM Global Profiles. The Global Profile enables PIM on the switch, and configures basic PIM parameters.


Creating a PIM Global Profile

Click on the Create icon **+**. Configure the profile as described below, then click on the **Create** button.

- **Profile Name** - User-configured profile name.
- **IPv4 Sparse Admin State** - Enables/Disables PIM-Sparse Mode (SM) protocol on the switch.
- **IPv4 PIM Bi Direction Status** - Enables/Disables Bi-Directional PIM on the switch.

Note: You can configure up to four (4) PIM Profiles; however two profiles cannot have the same values. For example, two profiles cannot be configured with both Sparse Mode and Bi-Directional status enabled, or with both disabled.

Editing a PIM Global Profile

Select the profile and click on the Edit icon  to bring up the Update PIM Global Configuration Screen. Edit the fields as described above then click on the **Update** button to save the changes to the server. The configuration will be applied and the status displayed on the Action Results Screen. Click the **Finish** button to return to the PIM Global Configuration Screen. Note that you cannot edit the Profile Name.

Note: Two profiles cannot have the same values. For example, two profiles cannot be configured with both Sparse Mode and Bi-Directional status enabled, or with both disabled.


Assigning a PIM Global Profile

Select a profile and click on the **Apply To Devices** button. Select an option (Use Switch Picker/Use Topology) and select the switch(es) to which you want to apply the profile and click the **Apply** button. The configuration will be applied and the status displayed on the Action Results Screen. Click the **Finish** button to return to the PIM Global Configuration Screen.

Removing a PIM Global Profile

To remove a profile from a switch(es), select the profile and click on the **Apply To Devices** button. Select the switches from which you want to remove the profile and click on the **Apply** button. The removal resets the values of Spare Admin State and Bi-Direction Status to "Disable".

Deleting a PIM Global Profile

To delete a Profile(s), select the Profile(s) in the table and click on the Delete icon , then click **OK** at the confirmation prompt. The configuration will be applied and the status displayed on the Action Results Screen. Click the **Finish** button to return to the PIM Global Configuration Screen.

Note: You can delete non-default profiles even the profile was assigned to switch. The deletion resets the values of Spare Admin State and Bi-Direction Status to "Disable". Also note that you cannot delete Default PIM profile, you can only remove the profile from the switches.


PIM Interface

The PIM Interface Screen is used to display information about configured PIM interfaces and to create or delete PIM interfaces. After enabling PIM on a switch by applying a PIM Global Profile, you must configure an IP interface as a PIM interface to enable multicast routing for VXLANs. An interface can be any IP router interface that has been assigned to an existing VLAN.


Displaying PIM Interfaces

You can view configured PIM interfaces by searching for PIM interfaces on a device(s) or by searching for specific PIM interface by name. Select a search option from the Search by drop-down menu (**Device** or **Interface**), then click on the **Select Devices** or **Select Interface** button to select the device(s)/interface(s) you want to view and click **OK**. If you select multiple devices, PIM interfaces common to those devices are displayed.


Creating a PIM Interface

Click on the Create icon . The Configure PIM Interface Screen appears. Select an option from the drop-down menu (Use Switch Picker/Use Topology) and click on the **Add/Remove Devices** button and select a device(s). If you select a single device, all IP interfaces configured on the device are displayed. If you select

multiple devices, only those interfaces common to all selected devices are displayed. Select an interface and click the **Create** button. The configuration will be applied and the status displayed on the Action Results Screen. Click **Finish** to return to the PIM Interface Screen.

Note: If there are no IP interfaces configured on a device, or no common interfaces among multiple devices, click on the Add  icon to bring up the VLANs application and configure the interface(s).

Deleting a PIM Interface

To delete an interface(s), select the interface(s) and click on the Delete icon , then click **OK** at the confirmation prompt.

PIM Candidate


The PIM Candidate Screen displays all configured PIM Candidate Profiles and is used to create, edit, and delete PIM Candidate Profiles. A PIM Candidate Profile is the Candidate Rendezvous Point (RP) Router and Candidate Bootstrap Router (BSR) configured on the switch. In PIM-SM, shared distribution trees are rooted at a common forwarding router, referred to as a Rendezvous Point (RP). The RP unencapsulates Register messages and forwards multicast packets natively down established distribution trees to receivers. The resulting topology is referred to as the RP Tree (RPT).

A Candidate RP Router is a PIM-enabled router that sends periodic Candidate RP advertisements to the Bootstrap Router (BSR). When a BSR receives a Candidate RP advertisement, the BSR may include the C-RP in its RP-set.


The role of a Candidate BSR is to keep routers in the network up to date on reachable Candidate RPs. The BSR's list of reachable Candidate RPs is also referred to as an RP set. There is only one BSR per PIM domain. This allows all PIM routers in the PIM domain to view the same RP set. A Candidate RP periodically sends out messages, known as C-RP advertisements. When a BSR receives one of these advertisements, the associated Candidate RP is considered reachable (if it has a valid route). The BSR then periodically sends its RP set to neighboring routers in the form of a Bootstrap message.

A Candidate BSR is a PIM-enabled router that is eligible for BSR status. To become a BSR, a Candidate BSR must become elected. A Candidate BSR sends Bootstrap messages to all neighboring routers. The messages include its IP address, which is used as an identifier, and its priority level. The Candidate BSR with the highest priority level is elected as the BSR by its neighboring routers. If two or more Candidate BSRs have the same priority value, the C-BSR with the highest IP address is elected as the BSR.


Creating a PIM Candidate Profile

Click on the Create icon . Select an option from the drop-down menu (Use Switch Picker/User Topology) and click on the **Browse** button to select the device on which you want to configure the PIM Candidate Profile. Configure the Candidate RP and Candidate BSR as described below.


Note: Devices will only be available/displayed if a PIM Interface has been configured on the device. If no PIM Interfaces have been configured, no devices will be available to create a PIM Candidate Profile.

- **Candidate RP**
 - **Candidate RP Address** - The IP address that will be advertised as a Candidate-RP. The IP address must belong to a PIM enabled interface. Only one RP address is supported per switch. Select a PIM Interface from the drop-down list. You can also click on the Add  icon to


go to the PIM Interface Screen and configure a PIM Interface.

- **Candidate RP Group Address/Prefix Length** - The group address for which the local router will advertise itself as a Candidate-RP and prefix length of the multicast group.
- **Candidate RP Bidir** - Enables/Disables Bi-Directional mode.
- **Candidate BSR**
 - **Candidate BSR Address** - The IP address of the Candidate BSR. Select a PIM Interface from the drop-down list. You can also click on the Add  icon to go to the PIM Interface Screen and configure a PIM Interface.

Editing a PIM Candidate Profile

Select the profile and click on the Edit icon  to bring up the Update PIM Candidate Screen. Edit the fields as described above then click on the **Update** button to save the changes to the server, and update the profile on the device(s).

Deleting a PIM Candidate Profile

To delete a profile(s), select the Profile(s) in the table and click on the Delete icon , then click **OK** at the confirmation prompt to remove the profile from the server and the device(s).

PIM Device View

The PIM Device View Screen is used to view PIM configurations on network switches. To view a configuration, select an option from the drop-down menu (Use Switch Picker/Use Topology), click on the **Browse** button, select a switch, and click **OK**. Click on a configuration setting (e.g., PIM Global, PIM Interface) to expand the view and see configuration details.

14.0 License

OmniVista 2500 NMS licensing is based on the license purchased. A user is allowed to manage up to the maximum number of devices allowed for that license. There are two licenses that can be purchased:

- **Node Management License** - Licenses all OmniVista applications except Virtual Machine
- **Manager (VMM) VMM License** - Licenses the VMM application.

The number of devices that can be managed with each license is determined by the License Key that the user is given and enters at installation. A new license can be imported and activated using the Add/Import Screen. Each License Screen (Node Management, VMM) also enables you to relicense/add a license. There are also licensing options that can be used to demo an application before purchasing a full Production License.

Note: If no Node Management License available in OmniVista, only the Add/Import Screen is displayed.

The License Management Screen, shown below, is displayed when you access the License application. The screen provides an overview of the number of devices being managed by the Node Management License (AOS Devices and Third-Party Devices) and the VMM License (VMs). The total number of devices that can be managed with the current license is shown at the top of each bar graph, and the number of devices currently being managed is displayed. The number of days remaining before each license expires is displayed at the bottom of each bar chart. If the number of devices currently being managed is more than 90% of the total number of devices that can be managed with the current license, the bar graph is displayed in Red, otherwise, it is displayed in Green. If the license is within 30 days of expiring, the license information at the bottom of the bar graph is displayed in Red; otherwise, it is displayed in Black.



Node Management License

The Node Management License (OV2500-NM) licenses all OmniVista applications except Virtual Machine Manager (VMM). OmniVista has been certified to manage up to 10,000 devices (includes AOS and Third-Party Devices). The Node Management License is activated using the Node Management

License Screen.

VMM License

The VMM License (OV2500-VMM) licenses the OmniVista Virtual Machine Manager (VMM) application. VMs can be deployed on vCenters, XenServers and Hyper-V Servers; and OmniVista supports a mixture of Hypervisor types. The OmniVista VM Manager application supports up to 5,000 VMs. More than 5,000 Virtual Machines are allowed, however a warning message will be displayed and an entry will be written to the VMM Log File. The VMM License is activated using the VMM License Screen.

Add/Import a New License

New Licenses are imported and activated using the Add/Import Screen. Note that you cannot downgrade a Node Management License or a VMM License. Also, make sure that the new license you are using will be adequate to manage the number of devices you were managing with the older version of OmniVista. If you are managing more devices than are allowed on your new license, you may lose some discovered devices.

OmniVista Licensing Options

There are three (3) types of OmniVista Licenses:

- **Starter Pack** - Is free and enables you to use OmniVista/OmniVista Applications (Node Management, VMM) on a limited basis.
- **Evaluation** - Gives you full use of OmniVista/OmniVista Applications (Node Management, VMM, Application Visibility), but for a limited time.
- **Production** - Gives you full use of OmniVista/OmniVista Applications without expiration.

Node Management License Options

	Starter Pack	Evaluation	Production
Device Count	20 (10 AOS, 10 Third Party)	Chosen at license generation (full OV functionality)	Chosen at license generation (full OV functionality)
Expires	No	60 Days	No

VMM License Options

	Starter Pack	Evaluation	Product
VMM Count	10	200	Chosen at license generation (full VMM functionality)
Expires	No	60 Days	No

License Management

The License Management Screen provides an overview of the number of devices/VMs being managed by each license type: Node Management License (AOS Devices and Third-Party Devices) and VMM License (VMs). The total number of devices/VMs that can be managed with the current license is shown at the top of each bar graph, and the number of devices/VMs currently being managed is displayed. The number of days remaining before each license expires is displayed at the bottom of each bar chart. If the number of devices/VMs currently being managed is more than 90% of the total that can be managed with the current license, the bar graph is displayed in Red, otherwise, it is displayed in Green. If the license is within 30 days

of expiring, the license information at the bottom of the bar graph is displayed in Red; otherwise, it is displayed in Black.

Node Management License

The License Node Management License Screen displays information about the current Node Management License and is used to relicense OmniVista.

Node Management License Information

The top part of the Node Management License Screen provides detailed information on the current OmniVista 2500 Node Management License. The lower portion of the screen displays usage information.

- **Expiration Date** - The date the current Node Management License expires.
- **License Type** - The current Node Management License type.
- **Product ID** - The Node Management License Product ID (OV2500-NM).
- **Unit** - The type of device managed by the Node Management License (ALE Devices, Third-Party Devices).
- **Max Count** - The maximum number of the devices (ALE and Third-Party) that can be managed with the current Node Management License.
- **Available** - The number of devices still available on the current Node Management License.
- **Usage** - The percentage of the maximum number of devices currently being managed. If the number of devices currently being managed exceeds 90 percent of the maximum number allowed on the license, the bar graph is displayed in Red; otherwise, it is displayed in Green.

Relicensing a Node Management License

To upgrade a Node Management License, or renew an expired license, click on the **Relicense** button to bring up the License Key Screen. Enter the new License Key and click **OK**.

VMM License

The License VMM License Screen displays information about the current VMM License and is used to relicense VMM. You can manage up to a total of 5,000 Virtual Machines (i.e., 5,000 VMs total on all Hypervisors).

VMM License Information

The top part of the VMM License Screen provides detailed information on the current OmniVista 2500 VMM License. The lower portion of the screen displays usage information.

- **Expiration Date** - The date the current VMM License expires.
- **License Type** - The current VMM License type.
- **Product ID** - The VMM License Product ID (OV2500-VMM).
- **Unit** - The type of VMs managed by the VMM License (VM).
- **Max Count** - The maximum number of VMs that can be managed with the current VMM License.
- **Available** - The number of VMs still available on the current VMM License.
- **Usage** - The percentage of the maximum number of VMs currently being managed. If the number of VMs currently being managed exceeds 90 percent of the maximum number allowed on the license, the bar graph is displayed in Red; otherwise, it is displayed in Green.

Relicensing a VMM License

To upgrade a VMM License, or renew an expired license, click on the **Relicense** button to bring up the License Key Screen. Enter the new License Key and click **OK**.

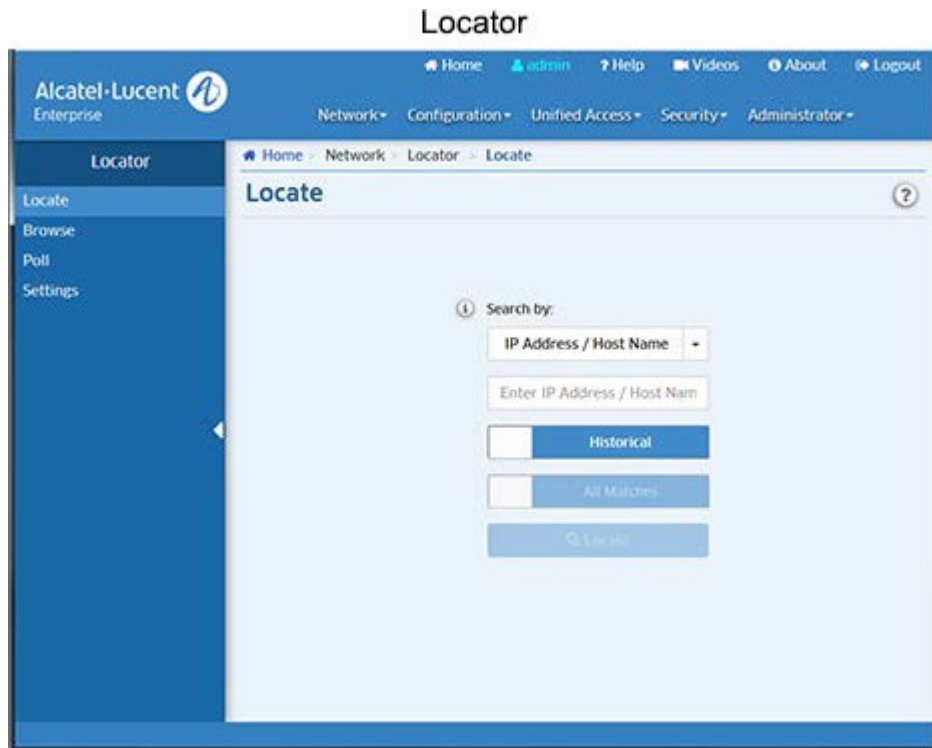
Add or Import License

The License Add or Import License Screen is used to activate a license, either by importing a license file, or entering a license key(s). To import a License File (.dat), download the file, then click on the **Download** icon, locate the file, import the file, then click on the **Submit** button. To activate a file using the license key, enter the license key in the License Key field. Note that you can enter multiple license keys. Each key must be entered on a new line (after entering a key, hit **Enter** to enter another license key). When you are done entering license keys, click on the **Submit** button.

15.0 Locator

The Locator application is a search tool within OmniVista. There are three (3) screens in the Locator application (accessed by clicking on the link on the left side of the screen) that are used to perform different functions:

- Locator's **Locator** screen enables you to locate the switch and slot/port that is directly connected to a user-specified end station. You can enter the end station's IP address or Host Name, MAC address, or Authenticated User ID to locate the switch and slot/port to which the end station is connected. Locator can perform a "Historical" search or a "Live" search. A "Historical" search is performed by searching a database of information that was previously established by polling network switches. A "Live" search, as its name implies, is performed by searching network switches in real time.
- Locator's **Browse** screen enables you to search in the "opposite direction" of the **Locator** screen. Instead of entering an end station's address to locate the switch and slot/port to which the end station is connected, the **Browse** screen enables you to search for and list all end stations connected to user-specified switch ports. The end stations are located by searching the Historical database. Locator cannot perform live searches from the **Browse** screen.
- Locator's **Poll** screen is used to immediately poll all of the discovered switches in the network for the latest information.
- Locator's **Settings** screen is used to set Locator timeout values and data retention policies.



Note: The Locator application supports IPv4 addresses only. IPv6 is not supported.

Locate Screen

To locate a switch, select a category from the Search by drop-down list, and enter the corresponding search criteria below (IP Address/Host Name, MAC Address, or Authenticated User ID). Choose to perform a Historical or Live search by clicking on the slider, then click **Locate**. Click here for more information on the Locator screen.

Although you can enter an end station's IP address, host name, MAC address, or Authenticated User ID to locate the switch and slot/port that is directly connected to the end station, Locator actually searches for the end station's MAC address. If you enter an IP address, host name, or Authenticated User ID, the first thing Locator does is find the corresponding MAC address. This MAC address is displayed in the search results, with a time stamp. The time stamp informs you how current the information is, which is especially important when performing historical searches. Locator then uses the MAC address to search for the switches, slots, and ports associated with the MAC. These are the final search results.

Whether performing a live search or a historical search, success in locating an end station depends on accurate topology information about switch-to-switch links. This information can be gathered using the Discovery application to discover new devices or re-discover existing devices; or by manually creating links using the Discovery - Link Screen.

Search Type

You can perform a Historical search or a Live search. As stated earlier, A "Historical" search is performed by searching a database of information that was previously established by polling network switches. A "Live" search, as its name implies, is performed by searching network switches in real time. The search process for each type is described below.

Historical Search

If a Historical search is performed, Locator first checks the list of Discovery Inventory List determine if the IP address, host name, or MAC address entered matches that of a known switch. If it does, a message is displayed informing the user and no further search is performed. If the IP address, host name, or MAC address entered does not to match a known switch, Locator assumes that the address is that of an end station and continues the search.

Live Search

If a Live search is performed, Locator will find all switches/slots/ports that meet the following criteria: the address entered was seen at the switch/slot/port, and the switch/slot/port is NOT connected to another switch device. If you select **1st Match Only**, only the first such switch/slot/port will be found. In most cases - as long as the network administrator has confidence in the consistency of the network's configuration - this result will be sufficient to locate the end station. If it is suspected that the first match may not be completely accurate, selecting **All Matches** will cause all switches/slots/ports that meet the criteria to be found and displayed.

Locator performs the search based on the search criteria entered, as described below.

For a Live Search, the Discovery Inventory List must contain switches that are supported by OmniVista (Alcatel-Lucent Enterprise and Third-Party). Every effort has also been made to support third-party devices, but that support is not guaranteed.

To successfully perform a live search for an IP address, the network's gateway device must be supported by OmniVista. Otherwise, Locator may not be able to resolve the IP address entered to a MAC address.

Locator searches for link information in the Topology database. This database must contain information about the links that exist between network switches. There are two methods of populating the Topology database with information about network links:

- Discover or re-discover devices using the Discovery application.
- Provide link information manually using the Discovery - Link Screen.

Search Results

When the search is complete, the search results are displayed in both an ARP Results Table and a Netforward Results Table. The ARP Results Table displays results matching the search criteria found in the ARP Table. The Netforward Table displays results matching the search criteria found in the Bridge Forwarding Tables. Click here for more information on search results.

Browse Screen

As stated earlier, the Browse screen enables you to search for and list all end stations connected to user-specified switch ports. The end stations are located by searching the Historical database. Locator cannot perform live searches from the Browse screen. To browse for a switch(es), click the **Select Switches** button to bring up a selection screen and select the switch(es) you want to discover. Click here for more information on the Browse screen.

Poll Screen

The Poll screen is used to update network information by immediately performing a poll of all discovered devices. Click here for more information on the Poll screen.

Locate

The Locator Locate Screen enables you to locate the switch and slot/port that is directly connected to a user-specified end station. You can enter the end station's IP Address, Host Name, MAC address, or Authenticated User ID to locate the switch and slot/port to which the end station is connected. Locator can perform a "Historical" search or a "Live" search. A "Historical" search is performed by searching a database of information that was previously established by polling network switches. A "Live" search, as its name implies, is performed by searching network switches in real time.

Locating a Switch

Select the type of search you want to perform and the search criteria as described below. The search results are described below.

- **Search by** - Select IP Address/Host Name, MAC Address, or Authenticated User ID from the drop-down list, then enter the search criteria.
- **Historical/Live Search** - Click on the slider to select a **Historical** Search or a **Live** Search. If you are performing a Live search, select **1st Match Only** to display only the first match found. In most cases - as long as the network administrator has confidence in the consistency of the network's configuration - this result will be sufficient to locate the end station. If it is suspected that the first match may not be completely accurate, select **All Matches** to display all switches/slots/ports that meet the selected criteria.

Search Results

When the search is complete, the search results are displayed in both the ARP Results Table and the Netforward Results Table. The ARP Results Table displays results matching the search criteria found in the ARP Table. The Netforward Table displays results matching the search criteria found in the Bridge Forwarding Tables.

You can also perform certain actions on specific devices/ports in the Netforward Results Table. Select a row in the table and click on the **Action** button at the top of the table and select one of the following options:

- **Locate On Map** - Launch the Topology application and display the selected device in Topology map view.
- **Quarantine Manager** - Launch the Quarantine Manager application for the selected MAC Address.
- **Port** - Update the port status in the table, or enable/disable the selected port.
- **Show ClearPass Authentication** - Launch the Authentication Records Screen in the Premium Services (BYOD) application for the selected MAC Address. This option is only available if a ClearPass Server is configured and connectivity can be established.
- **Show Access Guardian Diagnostics** - Launch the Diagnostics Screen in the Unified Profile application for the selected MAC Address.

Note: If the device you are searching for is a switch and not an end station, a notification will appear and you can click on the Locate on Map button to launch the Topology application and display the selected device in Topology map view.

ARP Results Table

In the ARP Results Table, Locator reports information for the end station you are searching on. Note that if you make changes to a switch's VLAN configuration, or if you make hardware changes (such as replacing a board), the results in the Initial Lookup area may not be correct. Before using Locator, re-poll such switches using the **Poll** Screen. This will ensure that the ARP tables are populated with current information. The ARP Results Table fields are defined below.

- **IP Address** - The IP address of the end station.
- **Devices IP Address** - The IP address of the device directly connected to the end station.
- **Device Name** - The name of the device directly connected to the end station.
- **MAC Address** - The MAC address of the device directly connected to the end station.
- **Timestamp** - The date and time the device was located.
- **End Station Name** - The name of the end station.
- **VPRN ID** - The VPRN ID of the device directly connected to the end station, if applicable. If multiple VRFs are configured on the device, the VRF ID is displayed. If none are configured (and if the feature is not available on the device), the column will display "Default", indicating that the switch is operating as a single routing instance.

Netforward Results Table

In the Netforward Results Table, Locator reports all switches/slots/ports that meet both of the following criteria: the address entered was seen at the switch/slot/port, and the switch/slot/port is NOT connected to another switch device. The Netforward Results Table fields are defined below. The table display will vary depending on the view option you choose - Location (default), Classification, Data Center, or Template, which is used to create custom views. If a ClearPass Server is configured and connected, a BYOD button will appear to enable you to view information on the ClearPass Server.

Location

- **MAC Address** - The MAC address of the end station connected to the selected device.
- **Devices IP Address** - The IP address of the device connected to the end station.
- **Device DNS Name**- The DNS Name of the device.
- **Device Name** - The user-configured device name.
- **Slot/Port** - The slot/port number on which the device was learned.
- **Port Alias** - The user-configured alias for the slot/port (configured on the device through the CLI).
- **Port Speed** - The port speed.

- **Port Admin Status** - The port administrative status (Up/Down).
- **Port Operational Status** - The administrative status of the port (Up/Down). A port is considered operational if the Admin Status is "Up" and the port is transmitting traffic.
- **Port Duplex Mode** - The port duplex mode (half duplex, full duplex, or auto duplex).
- **Timestamp** - The time the information was gathered.

Classification

- **MAC Address** - The MAC address of the end station connected to the selected device.
- **Auth User** - The 802.1x Authenticated User associated with the device connected to the end station, if applicable.
- **UNP** - The User Network Profile (UNP) that the device is associated with, if applicable.
- **ISID** - The ISID associated with the device.
- **Classification Source** - The Classification Policy by which the device was learned.

Data Center

- **MAC Address** - The MAC address of the end station connected to the selected device.
- **Service ID** - The Service ID associated with the device.
- **ISID** - The ISID associated with the device.
- **Classification Source** - The Classification Policy by which the device was learned.

Template

You can create an additional two (2) custom views by clicking on the **Custom Template** button, entering a **Template Name**, selecting the fields you want to display and clicking **OK**. You can also change the order of the fields when you are creating the template by dragging fields up or down in the list before clicking **OK**. The name of the new view will then be displayed in a button at the top of the Netforward Results Table, and can be used to view the selected fields. You can configure up to 2 new views. Creating an additional view will replace one of the previous views. Custom templates are associated with the current logged in user, so every user can have different custom templates. The available fields are defined below.

- **MAC Address** - The MAC address of the end station connected to the device.
- **Domain** - The Layer 2 domain: VLAN, VPLS, SPB, EVB within the switches where the MAC is found.
- **Slot/Port** - The slot/port number on which the device connected to the end station was learned. **Port**
- **Alias** - The user-configured alias for the slot/port (configured on the device through the CLI). **Port**
- **Speed** - The port speed of the device connected to the end station.
- **Port Duplex Mode** - The port duplex mode (half duplex, full duplex, or auto duplex).
- **End Station Name** - The name of the end station device.
- **Device DNS Name** - The DNS name of the end station connected to the device.
- **Device Name** - The user-configured switch system name of the device connected to the end station.
- **Last Updated** - The last time the information in the table was updated.
- **Auth User** - The 802.1x Authenticated User associated with the device connected to the end station, if applicable.
- **VRF ID** - The VRF ID of the device directly connected to the end station, if applicable. If multiple VRFs are configured on the device, the VRF ID is displayed. If none are configured (and if the feature is not available on the device), the column will display "Default", indicating that the switch is operating as a single routing instance.
- **UNP** - The User Network Profile (UNP) that the device is associated with, if applicable.
- **Classification Source** - The Classification Policy by which the device was learned.

- **Service ID**- The Service ID associated with the device.
- **ISID** - The ISID associated with the device.
- **Chassis** - The chassis number for devices supporting the Virtual Chassis feature.
- **Device IP Address** - The IP address of the end station connected to the device.
- **Time Stamp** - The time the information was gathered.
- **Port Operational State** - The port operational state (Up/Down).
- **VLAN ID**- The VLAN associated with the port.
- **Disposition** - The port disposition (e.g., Bridging/Filtering).

BYOD

Available only if ClearPass Policy Manager (CPPM) is defined in the BYOD application and connectivity can be established.

- **User Type** - The BYOD/ClearPass User Type authenticated through the device (e.g., Employee, Guest). You can also place the mouse over the user to view detailed ClearPass user information. BYOD is only supported on Alcatel-Lucent Enterprise Switches running AOS 6.4.6.R01 and later, AOS 6.6.5.R01 and later, AOS 7.3.4.R02 and later, and AOS 8.1.1.R01 and later. Note that the connection to the CPPM Server and Database must be configured properly in the Unified Access application to gather the necessary information for this field.
- **ClearPass Server** - The ClearPass Server name.
- **IP Address** - The IP address of the ClearPass Server.
- **User Name** - CPPM endpoint authentication information on the user who has logged in.
- **Category** - CPPM endpoint profiling information on the class of device (e.g., Computer, Smart Device, Access Points, VoIP Phone).
- **Family** - CPPM endpoint profiling information on the device family (e.g., Windows, Apple, Alcatel, Unix).
- **CP End Station Name** - CPPM endpoint profiling information on the device name (e.g., Windows Vista/7/2008, Apple iOS Device, Alcatel IP Phone, Wireless AP).
- **Host Name** - CPPM endpoint profiling information on the Host Name.
- **Sponsor Name** - CPPM guest information on the sponsor.
- **Visitor Name** - CPPM guest information on the visitor's name.
- **Visitor Company** -CPPM Guest information on the visitor's company.
- **Expires At** - CPPM guest information on the date when authorization will end.
- **Cert Valid From** - CPPM onboarding information of the date the certificate was issued.
- **Cert Valid To** - CPPM onboarding information of the date the certificate will expire.
- **MAC Vendor** - The manufacturer of the network equipment based on the Organization Unique Identifier (OUI).
- **ClearPass Time Stamp** - The time of the last authentication activity recorded in Clearpass.

Locate on Map

If the device you are searching for is a switch and not an end station, a notification will appear and you can click on the **Locate on Map** button to launch the Topology application and display a regional map in the Physical Network that contains the selected device. The device is automatically selected and centered in the map display.

Browse

The Locator Browse Screen enables you to search in the "opposite direction" of the **Locator** screen. Instead of entering an end station's address to locate the switch and slot/port to which the end station is connected, the Browse screen enables you to search for and list all end stations connected to devices. The end stations

are located by searching the Historical database. Locator cannot perform live searches from the Browse screen.

To browse for a switch(es), select an option from the drop-down menu (Use Switch Picker/User Topology) and click the **Select Device** button to select the switches you want to browse for, then click on the **Browse** button. The results will appear in the Netforward Results Table.

Browse Results

The browse results are displayed in the Netforward Results Table. The Netforward Results Table fields are defined below. The table display will vary depending on the view option you choose - Location (default), Classification, Data Center, Layer 3, or Template, which is used to create custom views. If a ClearPass Server is configured and connected, a BYOD button will appear to enable you to view information on the ClearPass Server.

You can also easily perform certain actions on specific devices/ports in the Netforward Results Table. Select a row in the table and click on the **Action** button at the top of the table and select one of the following options:

- **Locate On Map** - Launch the Topology application and display the selected device in Topology map view.
- **Quarantine Manager** - Launch the Quarantine Manager application for the selected device.
- **Port** - Update the port status in the table, or enable/disable the selected port.

Location

- **MAC Address** - The MAC address of the end station connected to the device.
- **Device IP Address** - The IP address of the switch to which the end station is connected.
- **Device DNS Name** - The DNS name of the end station connected to the device.
- **Device Name** - The user-configured switch system name of the device connected to the end station.
- **Slot/Port** - The slot/port number on which the device connected to the end station was learned.
- **Port Alias** - The user-configured alias for the slot/port (configured on the device through the CLI).
- **Port Speed** - The port speed of the device connected to the end station.
- **Port Admin Status** - The port administrative status (Up/Down).
- **Port Operational Status** - The administrative status of the port (Up/Down). A port is considered operational if the Admin Status is “Up” and the port is transmitting traffic.
- **Port Duplex Mode** - The port duplex mode (half duplex, full duplex, or auto duplex).
- **VLAN ID** - The VLAN associated with the port.
- **Disposition** - The port disposition (e.g., Bridging/Filtering).
- **Time Stamp** - The time the information was gathered.

Classification

- **MAC Address** - The MAC address of the end station connected to the device.
- **Auth User** - The 802.1x Authenticated User associated with the device connected to the end station, if applicable.
- **UNP** - The User Network Profile (UNP) that the device is associated with, if applicable.
- **VLAN ID** - The VLAN that is associated with the device.
- **Classification Source** - The Classification Policy by which the device was learned.

Data Center

- **MAC Address** - The MAC address of the end station connected to the device.
- **VLAN ID** - The VLAN associated with the device.
- **Service ID**- The Service ID associated with the device.
- **ISID** - The ISID associated with the device.

Layer 3

- **MAC Address** - The MAC address of the end station.
- **IP Address** - The IP address of the device connected to the end station.
- **VRF ID**- The VRF ID of the device directly connected to the end station, if applicable. If multiple VRFs are configured on the device, the VRF ID is displayed. If none are configured (and if the feature is not available on the device), the column will display "Default", indicating that the switch is operating as a single routing instance.

Template

You can create an additional two (2) custom views by clicking on the **Custom Template** button, entering a **Template Name**, selecting the fields you want to display and clicking **OK**. (You can change the order of the fields when you are creating the template by dragging a field up or down in the list before clicking **OK**.) The name of the new view will then be displayed in a button at the top of the Netforward Results Table, and can be used to view the selected fields. You can configure up to 2 new views. Creating an additional view will replace one of the previous views. Custom templates are associated with the current logged in user, so every user can have different custom templates. The available fields are defined below.

- **MAC Address** - The MAC address of the end station connected to the device.
- **Domain** - The Layer 2 domain: VLAN, VPLS, SPB, EVB within the switches where the MAC is found.
- **Slot/Port** - The slot/port number on which the device connected to the end station was learned.
- **Port Alias** - The user-configured alias for the slot/port (configured on the device through the CLI).
- **Port Speed** - The port speed of the device connected to the end station.
- **Port Duplex Mode** - The port duplex mode (half duplex, full duplex, or auto duplex).
- **End Station Name** - The name of the end station device.
- **Device DNS Name** - The DNS name of the end station connected to the device.
- **Device Name** - The user-configured switch system name of the device connected to the end station.
- **Last Updated** - The last time the information in the table was updated.
- **Auth User** - The 802.1x Authenticated User associated with the device connected to the end station, if applicable.
- **VRF ID**- The VRF ID of the device directly connected to the end station, if applicable. If multiple VRFs are configured on the device, the VRF ID is displayed. If none are configured (and if the feature is not available on the device), the column will display "Default", indicating that the switch is operating as a single routing instance.
- **UNP** - The User Network Profile (UNP) that the device is associated with, if applicable.
- **Classification Source** - The Classification Policy by which the device was learned.
- **Service ID**- The Service ID associated with the device.
- **ISID** - The ISID associated with the device.
- **Chassis** - The chassis number for devices supporting the Virtual Chassis feature.
- **Device IP Address** - The IP address of the end station connected to the device.
- **Time Stamp** - The time the information was gathered.
- **Port Operational State** - The port operational state (Up/Down).
- **VLAN ID**- The VLAN associated with the port.

- **Disposition** - The port disposition (e.g., Bridging/Filtering).

BYOD

Available only if ClearPass Policy Manager (CPPM) is defined in the BYOD application and connectivity can be established

- **User Type** - The BYOD/ClearPass User Type authenticated through the device (e.g., Employee, Guest). You can also place the mouse over the user to view detailed ClearPass user information. BYOD is only supported on Alcatel-Lucent Enterprise Switches running AOS 6.4.6.R01 and later, AOS 6.6.5.R01 and later, AOS 7.3.4.R02 and later, and AOS 8.1.1.R01 and later. Note that the connection to the CPPM Server and Database must be configured properly in the Unified Access application to gather the necessary information for this field.
- **ClearPass Server** - The ClearPass Server name.
- **IP Address** - The IP address of the ClearPass Server.
- **User Name** - CPPM endpoint authentication information on the user who has logged in.
- **Category** - CPPM endpoint profiling information on the class of device (e.g., Computer, Smart Device, Access Points, VoIP Phone)
- **Family** - CPPM endpoint profiling information on the device family (e.g., Windows, Apple, Alcatel, Unix).
- **CP End Station Name** - CPPM endpoint profiling information on the device name (e.g., Windows Vista/7/2008, Apple iOS Device, Alcatel IP Phone, Wireless AP)
- **Host Name** - CPPM endpoint profiling information on the Host Name.
- **Sponsor Name** - CPPM guest information on the sponsor.
- **Visitor Name** - CPPM guest information on the visitor's name.
- **Visitor Company** - CPPM Guest information on the visitor's company.
- **Expires At** - CPPM guest information on the date when authorization will end.
- **Cert Valid From** - CPPM onboarding information of the date the certificate was issued.
- **Cert Valid To** - CPPM onboarding information of the date the certificate will expire.

Poll

The Locator Poll Screen is used to immediately poll all of the discovered devices in the network for the latest information. OmniVista periodically polls the switches' Bridge Forwarding Tables to refresh the database used for historical searches. If you wish to force an immediate poll to refresh the database before your next search, click **Start Polling**.

This will begin an immediate poll of the Bridge Forwarding Tables of all known switches. This will refresh the historical data before you perform the next search. While polling is in progress, you can click **Stop Polling** to stop the polling process. When polling is complete, the Stop Polling button deactivates and the Start Polling button re-activates.

Note: Multiple OmniVista clients on the same OmniVista server cannot perform Locator polling simultaneously. Only a single OmniVista client can perform Locator polling at any one time. If a second client attempts to start a Locator poll while one is in progress, a "Cannot Start Polling" message will appear.

Settings

The Locator Settings Screen is used to set Locator timeout values and data retention policies. When you have configured the value(s), click the **Apply** button. The change takes effect immediately.

General

- **Historical Requests Response Timeout** - The amount of time, in seconds, that Locator will request historical data before timing out.
- **Live Requests Response Timeout** - The amount of time, in seconds, that Locator will request live data before timing out.
- **Locator Poll Requests Response Timeout** - The amount of time, in seconds, that Locator will poll a device before timing out.
- **802.1q Port Filtering** - This feature allows you to exclude 802.1q tagged ports from polling results and live searches when AMAP / XMAP is not operating or a link is not present on the tagged port. Filtering modes are described below.
 - **Standard Mode:** 802.1q Ports are included in polling and searches.
 - **Exclude Q-Tagged Ports:** 802.1q Ports are excluded from polling and searches.

Note: Virtual Machines communicate using tagged packets. If you are using VM Manager, 802.1q Port Filtering **must** be set to **Standard Mode** so that OmniVista can detect Virtual Machines with tagged frames.

Data Retention Policy

If Data Retention Policy is disabled, Locator will not remove data during polling and will accumulate unbounded data. To enable Data Retention Policy, move the slider to "Enabled" and set the retention period as described below.

- **Data Retention Period** - The number of days that Locator data will be retained (Default = 30).

Locator Data Statistics

When Data Retention is enabled, information that is older than the number of days specified in the Data Retention Period field is automatically deleted from the database. In addition, the number of days the data has been retained as well as the number of end station records being retained is listed. To refresh this information, click the Refresh icon. To purge all Locator data, click the **Purge All Locator Data** button, then click **Yes** at the confirmation prompt.

- **Retained Locator Data in Days** - The number of days Locator data has been retained.
- **Number of End Station Records Retained** - The number of end station records retained.

16.0 Notifications

The Notifications application is used to monitor traps and configure trap management tasks using the following screens:

- **Notifications Home** - Displays all traps received from network switches and provides basic trap information (e.g., severity level, date/time received). You can also use this screen to acknowledge, renounce, and clear traps, as well as poll devices for traps.
- **Trap Definition** - Displays a list of all supported traps, as defined in the MIBs, and provides a brief description of each trap. You can also edit a trap's severity level and trap synopsis.
- **Trap Responder** - Used to configure the response (if any) that you want OmniVista to take when a specified trap is received on the OmniVista server. The trap can be specified by severity level or through the use of filters. The response can take the form of an e-mail sent to a user-specified address and/or the execution of an external program or script on the OmniVista Server.
- **Trap Configure** - Used to configure traps for network devices.
- **Settings** - Used to set trap preferences (e.g., trap port, number of traps displayed).

Notifications


The screenshot displays the 'Notifications Home' screen. At the top, there is a navigation bar with 'Home', 'admin', 'Help', 'Videos', 'About', and 'Logout'. Below this is a breadcrumb trail: 'Home > Network > Notifications > Notifications Home'. The main content area is titled 'Notifications Home' and includes buttons for 'Acknowledge', 'Renounce', 'Poll for Traps', and 'Clear All'. There are also filters for 'All Maps & Devices', 'Severity: any', 'Acknowledged: false', and 'Most Recent'. A search bar is present, and the results are shown as a table with 12 rows. The table columns are Name, Severity, Acknowledged, Date Time, and Agent Name. The traps listed are all 'Normal' severity and 'False' acknowledged, with dates ranging from April 19, 2017, 3:06:48 pm to 3:19:00 am. The agent names include 'name' and 'ST10-CORE-01'.

Name	Severity	Acknowledged	Date Time	Agent Name
ala0vSwitchUp	Normal	False	Apr 19, 2017 3:06:48 pm	name
ala0vSwitchDown	Normal	False	Apr 19, 2017 3:01:39 pm	name
ala0vSwitchUp	Normal	False	Apr 19, 2017 2:16:23 pm	name
ala0vSwitchUp	Normal	False	Apr 19, 2017 2:16:23 pm	name
ala0vSwitchUp	Normal	False	Apr 19, 2017 1:44:32 pm	ST10-CORE-01
ala0vSwitchDown	Normal	False	Apr 19, 2017 1:36:26 pm	ST10-CORE-01
ala0vSwitchUp	Normal	False	Apr 19, 2017 6:36:00 am	ST10-CORE-01
ala0vSwitchDown	Normal	False	Apr 19, 2017 6:27:00 am	ST10-CORE-01
ala0vSwitchUp	Normal	False	Apr 19, 2017 4:17:52 am	ST10-CORE-01
ala0vSwitchDown	Normal	False	Apr 19, 2017 4:11:54 am	ST10-CORE-01
ala0vSwitchUp	Normal	False	Apr 19, 2017 3:22:29 am	ST10-CORE-01
ala0vSwitchDown	Normal	False	Apr 19, 2017 3:19:00 am	ST10-CORE-01


Notifications Home

The Notifications Home Screen displays configured traps received from network devices and provides basic trap information (e.g., severity level, date/time received). Click on a trap to display detailed trap information. You can use this screen to acknowledge, renounce, and delete traps, and manually poll devices for traps.

Viewing Traps

You can set the fields you want to display in the Notifications Table as well as the number of traps to display by clicking on the Settings button  at the top of the screen. You can also filter the display (e.g., display traps from selected devices, display traps by severity).

View Settings

Click on the Settings button  and configure the fields as described below. When you are finished, click **OK**.

- **Auto Refresh of Notifications Count** - The time, in minutes, to refresh the number of new notifications received. If new notifications are received during this time period, a message will appear at the top of the screen displaying the number of new notifications received, and giving you the option of immediately refreshing the Notifications Table display (Range = 1 - 15, Default = 2).
- **Auto Refresh of Notifications List** - The time, in minutes, to automatically refresh the Notifications Table display. The field must be enabled (On) to set this field (Range = 1 - 15, Default = 5).
- **Max No. of Notifications to Display** - The maximum number of notifications to display in the Notifications Table (Range = 1 - 10,000, Default = 1,000).
- **View Fields Selection** - Select the fields you want to display in the Notifications Table (e.g., Name, Agent Name, Synopsis). Click **OK** to save the changes. Note that the Severity, Acknowledged, and Date Time fields are always displayed.

Filters

There are several display filters available when viewing traps. You can view traps by network device, trap severity, trap acknowledgment state, and traps for a specific time period. Only those traps matching the filters you specify are displayed in the Notifications Table. This filters traps at the server level. Only those notifications matching the view criteria are pulled from the server and displayed in the Notifications Table.

Click on the **Filters** drop-down at the top of the Notifications Table and select one or more of the options below. When you are done selecting a filter(s), click on the "x" to close the Filter Window. The current filters are displayed next to the Filters drop-down.

- **View By** - Select which devices you want to view.
 - **All** - Select "All" to display configured traps from all network devices.
 - **Device** - Select "Device", then select Use Switch Picker or Use Topology then click on the **Browse** button to view traps from specific devices.
 - **Map** - Select "Map" and select a map to display traps from only the devices contained in the map.
- **Severity** - Select the Severity Level(s) of the traps you want to display: Critical, Major, Minor, Warning, Normal. You can select multiple Severity Levels by clicking on each level and then clicking anywhere outside the field box. De-select a previously-selected level by clicking on it. You can also select all levels by clicking on "Select All at the top of the drop-down menu). If you select "Remove All" levels, all traps are displayed, regardless of Severity Level ("0 Selected" - Default).
- **Acknowledged** - Select whether you want to view all traps (Any), Acknowledged Traps (True), or traps that have not been acknowledged (False).
- **Time Range** - By default, traps are displayed chronologically starting with the most recent trap. However, you can display traps for a specific time range to view older notifications. Select the **Custom** radio button and enter a beginning (From) and ending (To) date and time for the range. Select the **Most Recent** radio button to return to the default view. Note that up to 300,000 traps can be stored on the OmniVista Server. The number of traps stored on the server is configured on the Notifications Settings Screen.

Note: You can further filter the display by entering search criteria in the Search Field at the top of the table. The display will update to display only those entries that contain the search criteria. Delete the search criteria or click on the "x" at the far right of the Search Field to reset the display.

Basic Trap Information

- **Name** - The name of the trap as defined in the MIB.
- **Severity** - The severity level assigned to the trap (Normal/Warning/Minor/Major/Critical). Note that you can edit the severity level of a trap using the Trap Definition Screen.
- **Acknowledged** - Indicates whether or not the trap has been acknowledged (True) or not acknowledged (False). See Acknowledging/Renouncing/Deleting Traps for more information.
- **Date/Time** - The date and time the trap was received by the OmniVista server using the OmniVista server's system clock. However, for traps received that are "replays" of previously-generated traps, the date/time will be adjusted to the time that the original trap was sent. This is calculated by adjusting the time received by the difference between the current upTime of the source device and the upTime contained within the trap itself. Therefore, it is possible for new traps to be added to the display with old timestamps. So, if the network was down for hours, you may suddenly see traps appear from hours ago.
- **Agent Name** - The name of the device that generated the trap.
- **Synopsis** - A brief description of the trap.
- **Agent IP** - The IP address of the device that generated the trap.

Note: You cannot configure traps for wireless devices from OmniVista. However, you can configure traps on wireless devices and forward them to OmniVista for display.

Detailed Trap Information

The Detailed Information pane provides detailed information for the selected trap. If a field contains ellipsis (...), click on the field to display all of the information.


- **Name** - The name of the trap as defined in the MIB.
- **Synopsis** - A brief description of the trap. When a trap has variables associated with it, the values of some or all of the variables may appear in the synopsis. For example, in the trap synopsis "Link down on slot 6 port 2," the numbers "6" and "2" are trap variable values for the link down trap.
- **Agent IP** - The IP address of the device that generated the trap.
- **Agent Name** - The name of the device that generated the trap, if configured.
- **Date/Time** - The date and time the trap was received by the OmniVista server, using the OmniVista server's system clock. However, for traps received that are "replays" of previously-generated traps, the date/time will be adjusted to the time that the original trap was sent. This is calculated by adjusting the time received by the difference between the current upTime of the source device and the upTime contained within the trap itself. Therefore, it is possible for new traps to be added to the display with old timestamps. So, if the network was down for hours, you may suddenly see traps appear from hours ago.
- **Severity** - The severity level assigned to the trap in the Notifications Application's Trap Definitions Window (Normal/Warning/Minor/Major/Critical). Note that you can edit the severity level of a trap using the Trap Definition Screen.
- **Acknowledged** - Indicates whether or not the trap has been acknowledged (True) or not acknowledged (False). See Acknowledging/Renouncing/Deleting Traps for more information.
- **Description** - A detailed description of the trap as it appears in the MIB.
- **Up Time** - The length of time the device that sent the trap has been up (or the amount of time since the last reset), specified in days, hours, minutes, and seconds.

- **Source IP** - The IP address of the device that forwarded the trap to OmniVista.
- **Trap OID** - The trap object identifier number.

Note: You can click on the "Show More" link at bottom of the pane to display specific MIB variable information contained in the trap.

Acknowledging/Renouncing/Deleting Traps

You may want to temporarily remove some traps from the display by "acknowledging" the traps. Select a trap(s) and click on the **Acknowledge** button to remove the trap(s) from the display. Those traps will now only be displayed when you select "True" from the **Acknowledged** drop-down menu in the View Criteria area at the top of the screen. To return an "Acknowledged" Trap to the display, select "False" from the **Acknowledged** drop-down menu to display "Acknowledged" Traps, select the trap(s), and click on the **Renounce** button return the trap(s) to the display.

You can also delete an individual trap(s) by selecting the trap(s) and clicking on the Delete icon . To delete all traps from the table, click on the **Clear All** button at the top of the screen.

Note: Once traps are cleared or deleted, the traps are permanently removed from the OmniVista Server.

Polling Devices for Traps

To poll a device(s) for new traps click on the **Poll For Traps** button link at the top of the screen to bring up the Poll for Traps window. Select the device(s) you want to poll and click on the **Poll Now** button. The selected device(s) will immediately be polled for traps.

Trap Definition

The Notifications Trap Definition Screen displays a list of all the supported traps as defined in the MIBs, and provides basic trap information (e.g., severity level, date/time received). Click on a trap to display detailed trap information. You can also edit a trap or reset a trap to the installation defaults.

Viewing Trap Definitions

The Trap Definitions List displays basic trap information (e.g., name, severity level). Click on a trap to display detailed trap information. You can also search for traps by entering a search criteria (e.g., Name, Severity) in the "Search" field. Only those traps matching the search criteria are displayed. You can also filter the traps displayed, export the table to a .csv file, or print a copy of the table.

Basic Trap Information


- **Name** - The name of the trap as defined in the MIB.
- **Severity** - The severity level assigned to the trap in the Notifications Application's Trap Definitions Window (Normal/Warning/Minor/Major/Critical).
- **Synopsis** - A brief description of the trap.
- **Definition** - Indicates whether the trap is a "Default" trap or "Custom" trap.

Detailed Trap Information

The Detailed Information pane provides detailed information for the selected trap. If a field contains ellipsis (...), click on the field to display all of the information.

- **Name** - The name of the trap as defined in the MIB.
- **Trap OID** - The trap object identifier number.
- **Generic ID** - The Generic Trap ID number. Only SNMPv1 traps make use of a generic ID. For SNMPv2 and SNMPv3 traps, this field will show a value of zero.
- **Specific ID** - Trap specific ID number. Only SNMPv1 traps make use of a specific ID. For SNMPv2 and SNMPv3 traps, this field will show a value of zero.
- **Severity** - The severity level assigned to the trap (Normal/Warning/Minor/Major/Critical).
- **Synopsis** - A brief description of the trap. When a trap has variables associated with it, the values of some or all of the variables may appear in the synopsis. For example, in the trap synopsis "Link down on slot 6 port 2," the numbers "6" and "2" are trap variable values for the link down trap.
- **Agent IP** - The IP address of the device that generated the trap.
- **Definition** - The trap definition.
- **Description** - A detailed description of the trap as it appears in the MIB. Click on the "Show More" link to display a detailed MIB description.

Editing a Trap


You can edit a trap Severity Level or Synopsis. To edit a trap, select the trap in the Trap Definition List and click on the Edit icon . Edit the Severity Level and/or Synopsis and click on the **Apply** button. To return the field(s) to the default settings, select the trap in the Trap Definition List, click on the **Reset** button, then click **OK** at the confirmation prompt.

Trap Responder

The Notifications Trap Responder Screen displays all configured trap responders, and is used to create, edit, or delete Trap Responders. A Trap Responder enables you to specify a response (if any) that you want OmniVista to take when specified traps are received by OmniVista. You can specify the traps to which OmniVista will respond by IP address range and trap type. OmniVista can make the following responses to receipt of a specified trap:

- OmniVista can send an e-mail to any address you specify. You can specify the information you want included in the e-mail through the use of variables. Variables exist for information such as the trap name, synopsis, description, etc.
- Execute an external program or script on the OmniVista
- Server Forward traps to a specific IP address.

Creating a Trap Responder

Click on the Create icon  to open the Trap Responder Wizard. Complete the configuration as described below to configure the Agent, Trap Type, and Response. Only traps originating from the specified Agent and Trap Type will trigger the configured response. After completing all of the screens in the Wizard, click on the **Create** button.

Agent

The Agent is the IP address range for the Responder. The Responder will only respond to traps received from this IP address range. Enter an **Agent Start IP Address** and **Agent End IP Address** to specify the range. When you are finished, click the **Next** button to go the Trap Type window.

Trap Type

The Trap Type fields enable you to specify the traps to which OmniVista will respond by severity level; or, you can specify the traps to which OmniVista will respond using filters. Note that you can specify traps by severity level **or** filters, but you cannot specify traps using both severity levels and filters. (In other words, you cannot "AND" specified severity levels and specified filters to create an expression.) If you create a trap responder that specifies both severity level and filters, the trap responder will respond to all traps with the specified severity (even if they do not match the filter), and all traps that match the specified filters (even if they do not have the specified severity).

Severity Level

To select a severity level(s), move the corresponding slider to "Respond". OmniVista will send a response for any trap matching the selected severity level(s).

Filter

To use a filter, select "Any selected filter" (default) from the **Must Match** drop-down menu, then click on the **Filter** button to bring up the Filter Selection window and select a filter(s). OmniVista will send a response for any trap that matches **any** selected filter. If you select "All selected filters" from the **Must Match** drop-down menu, OmniVista will send a response for any trap that matches **all** of the selected filters.

To create a new filter, click on the Filter button to bring up the Filter Selection window, then click on the Add icon **+**. Enter a Filter Name and optional Filter Description, then click on the "Add new condition" link to add a filter condition. Click on the "Add new condition" link to add additional conditions, if necessary. When you are finished, click on the **Add** button to add the new filter. The following conditions can be configured for filters:

- Name Synopsis Agent
- Agent Name
- Date/Time
- Severity
- Acknowledge
- SNMP Variables

When you are finished configuring the Trap Type, click the **Next** button to go the Response window.

Response

Complete the fields as described below to configure the Response action to any traps matching the configured criteria.

- **Enable Responder** - Enable (On)/Disable (Off) the Responder.
- **Description** - Enter a description for the Responder.
- **Action** - Select the action that the Responder will take if the configured criteria is met, then configure the applicable fields as described below.

- **Send an E-Mail** - Responder will send an e-mail as configured below. Be sure to configure the E-Mail settings in the Preferences application (Preferences - System Settings - Email). OmniVista will not send an E-Mail Responder unless these settings have been configured.
 - **E-Mail To** - Enter an e-mail address(es). Multiple e-mail addresses must be separated by a semi-colon (e.g., john.smith@al-enterprise.com; jane.sullivan@al-enterprise.com).
 - **E-Mail Subject** - By default, the following is included in the E-Mail Subject Line: "OmniVista: Trap(s) Received \$TrapSeverityCount\$" (explained in the Trap Variables section below).
You can also enter any additional information you want to include in the subject line.
 - **E-Mail Body** - By default the variable \$Details\$ (explained in the Trap Variables section below) is included in the body of the e-mail. The variable automatically includes trap details in the body of the e-mail. You can also enter any additional information you want to include in the e-mail.
- **Run an Application on the Server**- Responder will run an application on the OmniVista Server as configured below.
 - **Command** - The command to be executed.
 - **Arguments** - The arguments to the command specified the Command field, or accept the default argument - the variable \$Synopsis\$ (explained in the Trap Variables section below).
 - **Start Directory** - The directory in which the command will be executed.
 - **Standard Input** - The standard input for the command in the **Standard Input** field, or accept the default standard input, the variable \$Details\$ (explained in the Trap Variables section below).
- **Forward Traps** -Responder will forward traps to the specified IP address .
 - **Destination IP** - The destination IP address. Only one IP address can be entered per Responder. However, you can create multiple Responders to forward the trap to multiple recipients.
 - **Destination Port** - The destination UDP port number (Default = 162).

When you are finished, click the **Next** button to go the Summary window.

Trap Variables

You can use the following variables when you configure an automatic trap responder. There are two types of variables: generic variables (which currently apply only to traps) and trap-specific variables.

Generic Variables

\$Details\$

For traps, this variable is equivalent to the following combination of text and trap-specific variables (trap- specific variables are described in the following section):

Trap Received:

\$TrapName\$ Severity:

\$TrapSeverity\$ Synopsis:

\$TrapSynopsis\$ Agent:

\$TrapAgent\$ Variables:

\$TrapVariables\$

Output Example:

Trap Received:

portPartitioned Severity:

Minor

Synopsis: Port jabber on slot 7 frtrunking port 1 instance 156 (port state alternated between enabled and disabled more than 50 times in 200 ms)

Agent: 128.251.30.27

\$Synopsis\$

For traps, this variable is equivalent to the trap-specific variable \$TrapSynopsis\$, which is a brief description of the trap.

Output Example: Port jabber on slot 7 frtrunking port 1 instance 156 (port state alternated between enabled and disabled more than 50 times in 200 ms)

*Trap Specific Variables***\$TrapName\$**

The name of the trap (as defined in the MIB)

Output Example: portPartitioned

\$TrapSynopsis\$

A brief description of the trap.

Output Example: Port jabber on slot 7 frtrunking port 1 instance 156 (port state alternated between enabled and disabled more than 50 times in 200 ms)

\$TrapDescription\$

A detailed description of the trap (as it appears in the MIB)

Output Example: A portPartitioned trap occurs when the physical port has transitioned through enable/disable states faster than 10 times in the past second...indicative of a flaky cable.

\$TrapSeverity\$

The severity level assigned to the trap in the Notifications application's Trap Definitions pane. The severity level can be Normal, Warning, Minor, Major, or Critical.

Output Example: Minor

\$TrapSeverityCount\$

A summary of the trap severity counts.

Output Example: (2 Critical)

\$TrapSeverityInt\$

The severity level assigned to the trap in the Notifications application's Trap Definitions pane, expressed as an integer. The severity level integer can be 1 (Normal), 2 (Warning), 3 (Minor), 4 (Major), or 5 (Critical).

Output Example: 3**\$TrapSnmpVersion\$**

The version of the trap request, either 1 (version 1) or 2 (version 2).

All traps sent with SNMP version 1 protocol are "version 1" trap requests. All traps sent with SNMP versions 2, 2c, or 3 protocol are "version 2" trap requests. There are actually two different types of trap requests (not three). The message packet in which trap requests are sent can be one of four different versions: 1, 2, 2c, or 3. When you use the AOS CLI to create a version 1 trap station, version 1 traps in version 1 protocol are sent to that station. When you use the AOS CLI to create a version 2 trap station, version 2 traps in version 2c protocol are sent to that station. When you use the AOS CLI to create a version 3 trap station, version 2 traps in version 3 protocol are sent to that station. The version 2 trap request itself is identical whether wrapped in a version 2 or version 3 packet.

Output Example: 1**\$TrapSource\$**

The IP address of the switch that generated the trap.

Output Example: 127.0.0.1**\$TrapUpTime\$**

The length of time the switch that sent the trap has been up (or the amount of time since the last reset).

Output Example: 21 hours, 35 minutes, 49 seconds**\$TrapAgent\$**

The IP address of the SNMP agent.

Output Example: 128.251.30.27**\$TrapV1Enterprise\$**

The enterprise name. This only applies to SNMP Version 1 traps.

Output Example: .1.3.6.1.4.1.800.3.1.1**\$TrapV1GenericID\$**

The generic trap number. This only applies to SNMP Version 1 traps.

¶Output Example: 6**\$TrapV1SpecificID\$**

The enterprise trap number. This only applies to SNMP Version 1 traps.

Output Example: 10**\$TrapVariables\$**

Describes all of the variables in the trap.

\$TrapVariable[1]\$, \$TrapVariable[2]\$...

Accesses the first (second, etc.) variable in the trap.


\$TrapVariable[someVariableName]\$

Accesses the trap variable by its name.


Summary

The Summary window displays the Responder configuration. Click on the **Create** button to create the Responder. If necessary, click on the **Back** button to make any changes before creating the Responder.

Editing a Trap Responder

Select the Responder in the Trap Responder List and click on the Edit icon . Edit the Agent and Trap Type as described above, then click on the **Apply** button.

Deleting a Trap Responder

Select the Responder(s) in the Trap Responder List, click on the Delete icon , then click **OK** at the confirmation prompt.

Viewing Trap Responders

The Trap Responder List displays basic trap responder information (e.g., trap type, response description). Click on a trap responder to display more detailed formation.

- **Agent** - The network region or IP address range configured for the Responder. The Responder will only respond to traps received from this region or IP address range.
- **Trap Type** - The severity level(s) or filter(s) configured for the Responder. The Responder will only respond to traps with this (these) severity level(s) or filter(s). You can also create filters for any or all of the Responder
- **Response Description** - The user-configured Response Description.
- **Enable Responder** - Indicates whether or not the Responder is enabled (True) or disabled (False).

Trap Configuration

The Notifications Trap Configuration Screen brings up the Trap Configuration Wizard, which is used to configure network devices to send traps to the OmniVista Server. Until you configure traps, you will not be able to receive or view any trap notifications on the Notifications Home Screen. Configuring traps is a two-part process: first you select the devices for which you want to configure traps, then you specify the traps you want to be sent to the OmniVista Server. At the end of the process, OmniVista displays a summary of all the switches and traps you selected, and indicates whether the configuration was successful or not.

Device Selection

The Device Selection window is used to select the devices for which you want to configure traps.

Server Information

- **IP Address (Read Only)** - The IP address of the OmniVista Server. This is automatically filled based on the OmniVista Server address you entered during installation.
- **Trap Port (Read Only)** - The destination trap port number on the OmniVista Server that receives alarms and traps. This field is automatically filled based on the Trap Port you configure on the Settings Screen.

Device Selection

- **Available Device Type** - Select the device type for which you want to configure traps. This will limit the available devices displayed for selection to the selected device type. Note that, no matter which type you choose, only devices that are "available" (Up) are displayed.
 - **All** - All available network devices are displayed.
 - **AOS** - Only AOS devices (pre-7x) are displayed.
 - **AOS 7x/8x** - Only AOS 7x/8x devices are displayed.
 - **6200** - Only OS6200 devices are displayed.
- **Available Devices** - Use the Switch Picker (Default) or the Topology application to select the specific devices for which you want to configure traps.

Note: You cannot configure traps for wireless devices from OmniVista, However, you can configure traps on wireless devices and forward them to OmniVista for display.

When you are finished, click the **Next** button to go the Configure Traps window. Click on the **Reset Trap Configuration** button to delete the configuration and start over.

Configure Traps

The Configure Traps window is used to select the traps you want to configure for the selected devices. The traps available for each device type are different. Depending on the device(s) selected, you will have the option to configure traps for each device type.

1. Click on a trap type to open the Trap Configuration pane for that device type (AOS, AOS 7x/8x, 6200).
2. Configure the trap information fields at the top of the pane:

- **Trap Subscription State**

On - The OmniVista Server will be notified about traps (Default).

Off - The OmniVista Server will not be notified about traps.

Delete - Deletes the trap configuration information previously saved to the switch. Use the Delete option when the OmniVista Server has been moved to a different computer and now has a different IP address.

- **Save**

All - Saves all trap configuration information specified, including port number, server IP address, selected switches, selected traps, state, and protocol.

Port Only - Saves only the port number and no other trap configuration information. Use this save option after configuring the OmniVista Server to receive traps on a different port. Note that the port number can be changed using the Settings Screen.

State Only - Saves only the state information (On, Off, or Delete) and no other trap configuration information. If Delete is selected, the entry for the OmniVista server is removed from the trap configuration table for the selected switches.

Traps Only - Saves only the trap information specified. Does not save Port or State information.

Protocol Only - Saves only the protocol used to send traps to the NMS server (SNMPv1, SNMPv2, or SNMPv3). **Applies only to AOS Switches.**

Note: If you select one of the "Only" Save options (Port Only, State Only, Traps Only, or Protocol Only), and no trap information had previously been configured for the specified

device, when you click Finish, the entire configuration will be saved. This is comparable to doing a Save "All."

- **Protocol** - Select the protocol used to send traps to the NMS server (SNMPv1, SNMPv2, or SNMPv3). (Default = SNMPv3) **Applies only to AOS/6200 Switches**).

3. Select the traps you want to enable in the **Select Traps to Enable** area.

4. Click on additional trap types and repeat the above steps to configure traps for additional device types. When you are finished, Click on the **Next** button to go to the Summary window. Click on the **Reset Trap Configuration** button to delete the configuration and return to the Devices Selection window and start over.

Summary

The Summary window displays the traps you configured each device type. Click on a device type (AOS, AOS 7x/8x, 6200) to review the configuration. If necessary, click on the **Back** button to return to the Configure Traps window and make any updates. When you are finished, click on the **Finish** button. Click on the **Reset Trap Configuration** button to delete the configuration and return to the Devices Selection window and start over.

Settings

The Notifications Settings Definition Screen is used to configure Notifications and Trap E-Mail settings. After making a change to one or more of the settings as described below, click on the **Apply** button.

Notification Configuration

- **Max No. of Notifications to Store at Server** - The maximum number of received traps that can be stored on the OmniVista Server. When a trap is received that exceeds the value in this field, the newly- received trap overwrites the oldest trap stored on the server. (Range = 1,000 - 300,000).
- **Trap Port Number** - The destination trap port number on the OmniVista Server that receives alarms and traps. The number entered in the Trap Port Number field must match the port number that the switch is configured to send traps to (Default = 162).
- **Use Trap Replay Polling** - If enabled (On), OmniVista will poll all discovered devices for missing traps at startup, and will continue to periodically poll devices for missing traps Default = On).
- **Generate OmniVista Switch Up/Down Traps** - If enabled (On), OmniVista will send Switch Up/Down traps (Default = On).
- **Use OmniVista Trap Absorption** - If enabled (On), similar traps received from non-AOS devices during the trap absorption period are "absorbed," and a 'trapAbsorbionTrap' trap is generated similar to existing AOS traps. This trap contains details, such as the total number of 'sufficiently-similar' traps received since the original trap. For example, if OmniVista receives a 'ChassisTrapsAlert' trap from a switch, OmniVista will 'absorb' all of the traps it receives from the same switch that are 'sufficiently similar' to 'ChassisTrapsAlert', until the trap absorption period expires. Note that two traps are considered to be "sufficiently similar" when their names, agent IP address, trap OID, severity, and enterprise OID (if defined) are same, and all their trap variables (if any) are also same. (Default = Off)
- **Absorption Period** - The amount of time, in seconds, OmniVista will "absorb" similar traps. OmniVista extends the trap absorption period when a 'trapAbsorbionTrap' trap is generated. For example, if a trap absorption period is set to 15 seconds and a 'sufficiently-similar' trap is received on the 8th second, the period for that trap is extended for another 15 seconds. If no 'sufficiently-similar' traps for a trap are received during the trap absorption period, the trap absorption period expires for

that trap. (Range = 0 - 600, Default = 15)

Trap E-Mail Configuration

The Trap E-Mail configuration fields are used to define the default values for OmniVista "trap responder" e-mails. Trap responder e-mails are e-mails that OmniVista generates automatically when specified traps are received from network devices. The Trap Responder Screen is used to configure OmniVista to send an e-mail when a specified trap is received. (Traps can be specified by severity level, or by use of a filter.) However, to prevent e-mail storms that would result from receipt of multiple traps, OmniVista does NOT send one e-mail per trap received. Rather, OmniVista "combines" responder e-mails to prevent storms. By default, OmniVista will send a "combined" responder e-mail when: one minute has passed since the first trap was received for which an e-mail was not generated, OR 100 traps have been received.

- **Maximum Trap Limit** - The number of received traps that will trigger a trap responder e-mail. Enter the number of received traps that will trigger a trap responder e-mail.
- **Maximum Time Limit** - The maximum time period, in seconds, that can elapse before a trap responder e-mail is triggered. The time period begins when a trap is received for which an e-mail was not generated.

17.0 PolicyView

The PolicyView application enables you to create Quality of Service (QoS) policies that specify QoS for network traffic. Policy rules are stored in a Lightweight Directory Access Protocol (LDAP) repository that is automatically installed with OmniVista and resides on the same device as the OmniVista Server. QoS-qualified devices in the network are notified when new or modified Policy rules are available on the LDAP repository via an SNMP interface. Software resident in the switch is responsible for retrieving the Policy rules from the LDAP repository, interpreting the Policy rules, and enforcing them on the switch.

When you first open the PolicyView application, links to the following options are displayed: Create Policies for Users and Groups, Create Policies for Resources, Create One Touch Policies, View/Modify Policies and Policy Lists, Expert Mode, and Create Policies for Applications.



The PolicyView application provides wizards to enable you to create specific QoS policy types (e.g., Application, Resource); and an "Expert" option that enables you to create more complex QoS Policies. These policies can be applied to all QoS-enabled devices in the list of All Discovered Devices or to selected QoS-enabled devices. These policies are created by associating a "Condition" with an "Action." A condition specifies criteria that, when true, will cause traffic to flow as specified by the associated action. A condition can specify criteria such as the following (a limited example):

- A source MAC address or a source IP address or a source VLAN ID, so that the condition applies to traffic originating from that source only
- A destination MAC address or a destination IP address or a destination VLAN ID, so that the condition applies to traffic flowing to that destination only.

An action specifies the treatment traffic is to receive when the criteria specified by the condition are true. This treatment may include the priority and bandwidth to be allocated to the traffic, its minimum and maximum output rates, and the manner in which packets are tagged upon egress from the switch (if at all).

The PolicyView application also provides a simplified "One Touch" mode that enables you to create QoS policies for data traffic and Access Control Lists (ACLs) with minimal effort and maximum simplicity. If you

use the One Touch option to create QoS policies for your network, there is no need to understand the underlying QoS definitions and constructs. The One Touch modes enable you to create QoS policies without bothering with the normal complexity associated with QoS. All QoS policies created using One Touch Policies are automatically applied to all QoS-enabled devices in the list of All Discovered Devices (Topology application).

The PolicyView application supports Provisioned QoS actions. By default, Provisioned QoS provides best-effort QoS in the switch. A Provisioned QoS action enables you to provide traffic with QoS other than best effort and to define the network resources, such as bandwidth and priority, to be made available to the traffic. When the criteria defined by the associated condition are true, traffic will be assigned to a queue that delivers the QoS specified by the action.

Important Note: Enabling Open Flow will consume all available TCAM resources. If Open Flow is enabled, you will be unable to configure QoS Policies. Any policies created before Open Flow is enabled will still function. However, you will be unable to create new policies.

Creating Policies for Users and Groups

The Users and Groups Policy is used to create/edit Unified Access Policies. Unified Policies are QoS Policies that can be applied to both wireline and wireless devices. For more information on creating Resources, [click here](#).

Creating Policies for Resources

The Resource Policy option is used to create/edit system resources for QoS Policies. Although you can use Policy View Expert Mode to create Policies for User Network Profiles (UNP), this can be time consuming. The Resource Policy option can be used to quickly create resources and resource groups that can be turned into Policies and added to Policy Lists. For more information on creating Resources, [click here](#).

Creating One Touch Policies

PolicyView provides a One Touch option that enables you to create One Touch Data and One Touch ACL Policies for traffic with minimal effort and maximum simplicity. One Touch Data policies enable you to assign a desired quality of service - Platinum, Gold, Silver, or Bronze - to all traffic flowing to, and originating from, specific data servers. One Touch ACL Policies enable you to create ACL Policies to all traffic flowing to, and originating from, specific Network Groups. For more information on One Touch Policies, [click here](#).

View/Modify Policies and Policy Lists

This option enables you to view and modify all Policies and Policy Lists stored in the LDAP Server. To view the policies, click on the **Select Devices** button to open the Device Selection Wizard and select the switches you want to view. The devices will appear in the Selected Devices Table. Select a device in the table to display the Policies and Policy Rules for the selected device. For more information, [click here](#).

Expert Mode

In the **Expert** mode, conditions and actions are not created automatically; and the user defines the devices to which the policies are assigned. The Expert mode enables you to create conditions and actions manually, by specifying each individual parameter. In the Expert mode, you can create conditions that specify MAC addresses, IP address, protocols, VLAN IDs, specific DSCP or TOS values, or specific 802.1 priority values. For more information on creating policies on the Expert mode, [click here](#).

Creating Policies for Applications

This option enables you to create Application Visibility Policies and Policy Lists for Application traffic flows. For more information on Application Visibility Policies, [click here](#).



QoS-Qualified Devices

A QoS-qualified device is a device that can support the PolicyView application and provisioned QoS. AOS devices are qualified devices. QoS-qualified devices are identified during the discovery process. The list of QoS-qualified is available and can be displayed on Expert Mode Screen.

Saving Changes to the Switch

When PolicyView is executed, it writes the address of the LDAP repository to each QoS-qualified switch in the Inventory List in the Discovery application. The LDAP address is written to the running configuration of the switch. For this reason, once PolicyView has executed, all switches are left with their running configuration in the "Unsaved" state (indicating that the running configuration has changes that have not been saved to the working directory). When a switch reboots, its running configuration is lost, so it is important to save the running configuration, and then to save the running configuration to the certified directory after PolicyView has executed. To do this, follow the steps below.

Note: All changes made to the switch configuration will be saved, including any changes made via the CLI, WebView, or other OmniVista applications, in addition to the changes made by the PolicyView application.

1. Go to the Discovery application to view all discovered devices in the Inventory List.
2. Scroll right to the "Changes" column and sort the list according to the switch configuration state.
3. Select all switches with "Unsaved" changes. Click on the Operations  at the top of the list and select **Save to Running**. The "Changes" field will display "Uncertified" when the changes are saved to the Running directory.
4. Select all switches with "Uncertified" changes. Click on the Operations  at the top of the list and select **Copy Working/Running to Certified**. The "Changes" column will go blank when the Working/Running Directory is saved to the Certified Directory (this may take a few minutes).

Note: You could also perform the operation above using the operations in the Topology application.

Required Traps

You must configure the switches in the network to send OmniVista the traps that are needed by the PolicyView application. To configure traps for one or more devices, go to the Topology application, select the device(s) and select **Notifications - Configure Traps** from the Operations panel. The Trap Configuration Wizard appears with the selected switches. PolicyView requires the following traps:

- 8 - policyEventNotification

Note: See the Topology application help for step-by-step instructions for configuring traps.

Policy Precedence and Conflicts

PolicyView enables you to define the precedence of policies created in PolicyView. A policy rule's precedence determines which policy will take effect in the rare case of a conflict. QoS policies can be created through the CLI, through WebView, and through SNMP MIB browsers as well as through PolicyView. Policies created through the CLI, WebView, or MIB browsers are not written to the LDAP repository and are not manageable through the PolicyView application.

Note: It is highly recommended that network administrators who use PolicyView to create policies do NOT use any outside management tools for creating policies, conditions, or actions.

Policies created in PolicyView are assigned a precedence value between 30001-65535. However, precedence values 30001-65535 are not reserved for PolicyView policies. Policies can also be created using the CLI, WebView, or a MIB browser, and these policies can be assigned any precedence value between 0- 65535. Therefore, it is possible to assign these policies the same precedence that is assigned to policies created through the PolicyView application. For this reason, if you are creating policies using PolicyView as well as outside management tools (which is NOT recommended), do not assign precedence values between 30001-65535 to any policies created outside of the PolicyView application.

- One Touch Voice policies have precedence values between 45000 and 65535.
- One Touch Data policies have precedence values between 40000 and 44999.
- Expert Mode policies have precedence values between 30000 and 39999.

Creating Policies for Users and Groups


The Users and Groups Policy is used to create/edit Unified Access Policies. Unified Policies are QoS Policies that can be applied to both wireline and wireless devices.

Unified Policies

The PolicyView Unified Policies Screen application displays all configured Unified Policies and is used to create, edit, delete, and view Unified Policies. Unified Policies are QoS Policies that can be applied to both wireline and wireless devices. Unified Policies are created using a wizard that guides you through each of the steps needed to create the Policy and apply the Policy to devices in the network.

Note: Unified Policies are only displayed in the Unified Policies Table. They are **not** displayed with other configured QoS Policies in the Expert Mode Existing Policies Table.

Creating Unified Policy

Unified Policies are created using a wizard that guides you through each of the steps needed to create the policy and apply the policy to devices in the network. To create a Unified Policy, click on the Create icon . The wizard will then guide you through the following screens:

- Configuration - Basic policy configuration (e.g., Policy Name, Precedence)
- Device Selection - Specify the devices to which you will apply the policy
- Set Condition - Specify the conditions that must be true before traffic will be allowed to flow. Set Action
- - Specify parameters for the traffic that will flow.
- Validity Period - Specify the time period for the policy to be effective. Review - Review the policy details before creating the policy.

Note: As you configure a policy, conditions and actions are verified against the devices selected for the policy. If a condition or action is not supported by one of the selected devices, an error message will appear indicating the error and corrective action to be taken.


Applying a Unified Policy to the Network

After configuring and saving a policy(ies), you must apply the policy(ies) by notifying the switches in the network. When you click on the Notify All button, all of the policies listed in the Existing Unified Policies Table are applied to all of the devices configured for each policy. To apply the policy(ies) only to certain devices, select an option from the drop-down menu (Use Switch Picker/Use Topology), click on the Select Device button and select the device(s); then click the Notify Selected button.

After notifying the devices, you can view the status of the re-cache operation, by clicking on the Status button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the Audit application, including an indication of any error that may have occurred. Note that the re-cache operation for each switch occurs in a separate thread and may take some time. Any errors that occur will also be reported in the server.txt file, in the Audit application.


Note: When you notify network switches, all QoS-enabled switches flush their policy tables and reload policies from the LDAP repository, which is very expensive in terms of switch resources and time. It is recommended that you review all policies that you have created and apply them at the same time to minimize switch downtime.

Editing a Unified Policy

To edit a policy, select the policy in the Existing Unified Policies Table and click on the Edit icon . Use the wizard to make any edits. When you are done, apply the edited policy to the network.

Note that if you modify a policy and select different device types at the Device Selection Step (AOS/Wireless), a warning dialog will be displayed if the condition in the policy is not supported on one of the selected device types. For example: “Condition mis-match: condition (L2 MACs and L4 Service) is not valid for selected device. Do you want to remove the mis-match conditions?” If you select Yes, the mis-matched conditions will be removed from edited policy. Otherwise, the newly selected devices will be removed from Device Selection list.

Deleting Unified Policy

To delete a policy(ies), select the policy(ies) in the Existing Unified Policies Table and click on the Delete icon  then click **OK** at the confirmation prompt.

Policy Information

The Existing Unified Policies Table displays information for all configured Policies. You can also click on a policy to view detailed information about the Policy (e.g., Condition, Action).

- **Policy Name** - The name of the Policy.
- **Scope** - The scope of the Policy (e.g., IP Filtering, Provisioned).
- **Precedence** - The Precedence value of the Policy (0 - 65535).
- **Status** - Indicates whether or not the Policy has been saved to the LDAP Server.
- **Enable** - Indicates whether or not the Policy is enabled.
- **Save** - Indicates whether or not the rule will be recorded during a snapshot command.
- **Log Matches** - Indicates whether or not matches to this rule are logged in the QoS Log.

- **Reflexive** - Indicates whether or not the Policy is reflexive. Reflexive Policy Rules allow specific return connections that would normally be denied (Yes/No).
- **Default List**- Indicates whether or not the Policy is saved to the Default Policy List. By default, a Policy Rule is added to this list when it is created. A Policy Rule remains a member of the Default List even when it is subsequently assigned to additional Policy Lists.
- **SLA Policy Trap** - Indicates whether or not an SLA Policy Trap is configured for the policy.

Config for Policy

The Unified Policies Config for Policy Screen is used to configure basic Policy parameters. When you have completed all of the parameters, click the **Next** button at the bottom of the screen or click on Device Selection on the left side of the screen to move to the next step.

- **Name** - The Policy name.
- **Precedence** - The Policy precedence. By default, the precedence field is pre-filled with the lowest unused precedence value (Range = 0 - 65535).

Click on **Show Advanced Options** to display and configure the options below:

- **Default List** - Adds the rule to the QoS Default Policy List. Default is **No**
- **No Enabled** - Enables the policy. Default is **Yes**
- **Save** - Marks the policy rule so that it may be captured as part of the switch configuration. Default is **Yes**
- **Log Matches** - Configures the switch to log messages about specific flows coming into the switch that match this policy rule. Default is **Yes**
- **Send Trap** - Enables traps for the Policy. Default is **No**
- **Reflexive** - Enables support for the Reflexive for the policy. Reflexive policies allow specific return connections that would normally be denied. Default is **Ignore**.

Note: The Config for Policy Screen for Unified Policies is similar to Config for Policy Screen in Expert mode. However, Unified Policies created for Wireless Controllers will accept the "No Reflexive" option.

Device Selection

The Unified Policies Device Selection Screen is used to select the switches to which you want to apply the Policy. Select an option (Use Switch Picker/Use Topology), and select the device(s). Click on the **Next** button at the bottom of the screen or click on Set Condition on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen and add/delete devices.

Note: In Expert Mode, you can only select AOS Devices for Policy creation. However, you can select wireline and wireless devices when creating Unified Policies. Also note that you cannot select IAP Devices when creating Unified Policies.

Set Condition

The Unified Policies Set Condition Screen contains a list of Conditions that you can configure for the Policy (e.g., MAC Condition, IP Condition). When you create a Condition, the Condition(s) you configure must be true before traffic is allowed to flow. Click on a Condition to display the configuration options for the Condition. (Click again on the Condition to close the configuration options.) When you have completed all of the parameters for the Condition(s), click the **Next** button at the bottom of the screen or click on Set Action on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.


A brief description of each Condition is provided below. Click the hyperlink for each Condition for detailed configuration instructions.

- **L2 MACs** - Create a Condition that applies the policy to traffic originating from a MAC address/group/range or to traffic flowing to a MAC address/group. (Note that for Wireless Controllers, MAC Addresses cannot contain wildcard characters).
- **L3 IPs** - Create a Condition that applies the policy to traffic originating from an IP address/network group or to traffic flowing to an IP address/network group. (Note that any IP address can be masked).
- **L3 DSCP/TOS** - Create a Condition that applies the policy to traffic with a specified value in either the DSCP (Differentiated Services Code Point) byte or in the IP TOS (IP Type of Service) byte. Both DSCP and IP TOS are mechanisms used to convey QoS information in the IP header of frames.
- **L4 Services** - Create a Condition that applies the policy to traffic flowing between two TCP or UDP ports, or to all traffic originating from a TCP or UDP port, or to all traffic flowing to a TCP or UDP port. You can also create a Condition using an existing service/service group.

Note: AOS Devices support most of above Conditions. However, Wireless Controllers do not. Please refer to detailed notes of each condition below for supported conditions.

L2 MACs

A MAC Condition applies the Policy to traffic flowing from/to a MAC Address/Group. Note that Layer 2 Conditions (conditions that specify MAC Addresses) are "lost" when traffic passes through a router. For this reason, it may be advisable to specify other types of Conditions (such as a Layer 3 Condition, which specifies IP Addresses) when traffic is expected to travel more than one router hop.

Select the parameter(s) you want to configure by selecting the applicable checkbox. Click on Single to configure a single MAC Address or Group to configure a MAC Group, then enter a MAC address or select a MAC Group from the drop-down menu. (You can also click the Add icon  to go to the Groups application and create a new MAC Group.)

- **Source MAC Address/MAC Group** - Configuring a Source MAC Address/Group Condition restricts the policy to traffic that flows from this MAC Address/Group only. If you do not select this option, you are effectively stating that the Source MAC Address/Group traffic is not a criterion for the policy.
- **Destination MAC Address/MAC Group** - Configuring a Destination MAC Address/Group Condition restricts the policy to traffic that flows to this MAC Address/Group only. If you do not select this option, you are effectively stating that the Destination MAC Address/Group traffic is not a criterion for the policy.
- **Source MAC Range** - Configuring a Source MAC Range Condition restricts the policy to traffic that flows from this MAC Range only. If you do not select this option, you are effectively stating that the Source MAC Range traffic is not a criterion for the policy.


Notes:

- Conditions that specify both a source and a destination MAC address may be rejected by some switch platforms as invalid. However, if you wish to create policies for both source and destination traffic, you can create one policy for the source traffic and a second policy for the destination traffic.
- MAC addresses may contain the wildcard character *. However, one * character must be entered for each individual hex digit in the MAC address: for example, **00435C:*******, not **00435C:***.
- The following MAC address ranges are assigned to Alcatel-Lucent Enterprise voice devices and Alcatel-Lucent Enterprise IP phones. You can create Conditions specifying these address ranges using the MAC Address tab.

- Voice Devices
 - 00809F3A0000 - 00809F3AFFFF
 - 00809F3B0000 - 00809F3BFFFF
 - 00809F3C0000 - 00809F3CFFFF
- IP Phones
 - 00809F3D0000 - 00809F3DFFFF
- Multi-Media Devices
 - 00809F3E0000 - 00809F3EFFFF
 - 00809F3F0000 - 00809F3FFFFFF
- Source MAC Range is not supported on AOS Devices.
- Source MAC Group and Destination MAC Address/MAC Group are not supported on Wireless Controllers.
- MAC Conditions are not supported on IAP Devices.

L3 IPs

An IP Condition applies the Policy to traffic originating from, or flowing to, an IP Address/Network group. Any IP Address can be masked. Note that a Condition that specifies both a Source and Destination IP Address/Network Group will be rejected by the switch as invalid. However, if you wish to create policies for both Source and Destination traffic, you can create one policy for the Source traffic and a second policy for the Destination traffic.

Select the parameter(s) you want to configure by selecting the applicable checkbox. For Source/Destination IP Address, click on Single to configure a single IP Address (and Shorthand or Subnet Mask, if applicable), or click on Group to configure a Network Group, then enter an IP Address or select a Network Group from the drop-down menu. (You can also click the Add icon  to go to the Groups application and create a new Network Group.)

- **Fragment (not available for Wireless controllers)** - Select this checkbox to restrict the policy to TCP packet fragments.
- **Source IP Address/Network Group** - Configuring a Source IP Address/Network Group Condition restricts the policy to traffic that flows from this IP Address or Subnet Mask/Network Group only. If you do not select this option, you are effectively stating that the Source IP Address or Subnet Mask/Network Group traffic is not a criterion for the policy.
- **Destination IP Address/Network Group** - Configuring a Destination IP Address/Network Group Condition restricts the policy to traffic that flows to this IP Address/Network Group only. If you do not select this option, you are effectively stating that the Destination IP Address or Subnet Mask/Network Group traffic is not a criterion for the policy.
- **Multicast IP Address Range (not available for Wireless Controllers)** - Configuring a Multicast IP Address/Group Condition restricts the policy to traffic that flows to this IP Multicast Address Group only. If you do not select this option, you are effectively stating that the Destination IP Multicast Address or Subnet Mask/Group traffic is not a criterion for the policy.

Notes:

- When configuring an IP Address Condition, you can also click either the **Shorthand Mask** or **Subnet Mask** button to configure a Subnet Mask. If you are using a Shorthand Mask, select a value from the Shorthand Mask drop-down list. If you are using a full Subnet Mask, enter the mask in the IP Subnet Mask field. Note that the * wildcard character is not allowed in IP addresses.

- Short hand Mask and Group are ignored when applying Unified Policies to Wireless Controllers.
- Source Group, Destination Group and Multicast are not supported on Wireless Controllers.

Important Note: When creating an IP Condition for a **NAT** Action you must specify a Network Group in the Condition. NAT will only work when both the Condition and Action specify network groups. To create a "One-to-Many" Condition and action, create a Network Group with a single entry for the Condition.

L3 DSCP/TOS

A DSCP/TOS Condition applies the Policy to incoming traffic that has a specified value in either the DSCP (Differentiated Services Code Point) byte or in the TOS (Type of Service) byte. Both DSCP and TOS are mechanisms used to convey QoS information in the IP header of frames. DSCP and TOS are mutually exclusive - you can use either DSCP or TOS but not both. Click on the applicable button (DSCP or TOS) and enter a value.




- **DSCP** - Defines the QoS treatment a frame is to receive from each network device. This is referred to as per-hop behavior. If you are using DSCP, you can define any value in the range 0 - 63 as the DSCP value in the IP header of the frame. Traffic that contains this value will match this condition.
- **TOS** - A TOS value creates a condition that applies the policy to traffic that has the specified TOS value in the IP header of frames. Enter any value from 0 - 7 to specify the value of the precedence field in the TOS byte that will match this condition. A value of 7 has the highest precedence and a value of 0 has the lowest.

Notes:

- Please refer to the Switch Release Notes for information on the specific QoS functions available on various current platforms and combinations of hardware/firmware.
- You cannot create a policy condition based on DSCP or TOS values for Wireless Controllers/IAPs. DSCP/TOS conditions are ignored when applying Unified Policies to Wireless Controllers/IAPs.

L4 Services

A Service Condition applies the policy to Service Protocol traffic (TCP or UDP) flowing from/to two TCP or UDP ports, or to traffic flowing from/to a TCP or UDP Service or Service Group. Select a type of Service Condition you want to configure, then configure the parameter(s) as described below.

- **Protocol Only** - Select **TCP** or **UDP** to create a condition for a Service Protocol only.
- **Port(s)** - To configure the Condition for a specific Service Port, select a **Source** and **Destination** Port from the drop-down menu to specify a specific port for the service you selected. You can also click on the Add icon  to go to the Groups application and create new Service Ports.
- **Service** - Select a Service from the drop-down menu. You can also click on the Add icon  to go to the Groups application and create a new Service.
- **Service Group** - Select a Service Group from the drop-down menu. You can also click on the Add icon  to go to the Groups application and create a new Service Group.

Notes:

- Wireless Controllers do not have source and destination ports. They only contain a unique service port. Therefore, you cannot specify both Source and Destination port for Wireless Controllers.

Set Action

The Unified Policies Set Action Screen contains a list of Actions that you can configure for the Unified Policy (e.g., QoS, NAT). A Policy Action enables you to specify the treatment traffic is to receive when it flows. This includes the priority the traffic will receive, its minimum and maximum output rates, and the values to which specified bits in the frame headers will be set upon egress from the switch. When the Conditions specified by the Policy Condition are true, traffic will flow as specified by the Policy Action.

Click on an Action to display the configuration options for the Action. (Click again on the Action to close the Action.) When you have completed all of the parameters for the Action(s), click the **Next** button at the bottom of the screen or click on Validity Period on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.

Actions

A brief description of each Action is provided below. Click the hyperlink for each Action for detailed configuration instructions.

- **QoS** - Create an Action to specify QoS actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action. Quality of Service applies to Session Type for wireless devices. Quality of Service is not supported on IAP devices and is ignored when applied to those devices.
- **TCM** - Create an Action to specify Tri-Color Marking (TCM) actions to impose on traffic that meets the configured policy condition(s). TCM provides a mechanism for policing network traffic by limiting the rate at which traffic is sent or received on a switch interface. TCM meters traffic based on user-configured packet rates and burst sizes and "marks" the metered packets as green, yellow, or red based on whether the traffic meets the configured rates. This "color marking" determines the packet's precedence when congestion occurs. TCM is not supported on wireless devices and is ignored when applied to those devices.

QoS

The QoS Policy Action option enables you to specify QoS actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

- **Disposition** - Set the Action to **Accept** or **Drop** traffic that meets the configured condition(s).
- **Quality of Service (QoS) Parameters** - Specify the QoS priority the traffic will receive if it meets the configured condition(s).
 - **Platinum** priority provides the highest quality of service (and maps to a firmware priority of 7).
 - **Gold** provides the next-highest quality of service (and maps to a firmware priority of 5).
 - **Silver** provides the next-highest quality of service (and maps to a firmware priority of 3).
 - **Bronze** provides the same quality of service as best effort (and maps to a firmware priority of 1). A separate egress queue is maintained in the hardware for traffic of each different priority.
- **Output Flow Setting** (not supported on IAP Devices and is ignored when applied to those devices)
 - **Max Output Rate (kbits/sec)** - Specify the maximum amount of traffic, in kilobits-per-second, which is guaranteed to be transmitted from the port. Even if no other traffic exists, the output will be limited to the rate specified here.
 - **Set Color of Packet** - Enables/Disables Three Color Marking (TCM) for output traffic flows. This parameter is not supported on wireless devices and is ignored when applied to those devices.

- **Output Mapping** (not supported on IAP Devices and is ignored when applied to those devices)
 - **802.1p Priority Level** - If you want outgoing packets tagged with an 802.1p priority level, set the **802.1p Priority Level** field to any value between 0 to 7 to specify the desired outgoing 802.1p priority for the traffic. A value of **7** indicates the highest priority and a value of 0 indicates the lowest priority. Note that for ports that are configured for 802.1q, this value is used in the 802.1q header and indicates the outgoing priority of the frame. When a frame is de-queued for transmission, it is assigned the priority of the queue and mapped to the outgoing 802.1p priority. This priority is combined with the VLAN group ID to create the 802.1p/q header for transmission. Note that if traffic matches the criteria specified by the policy condition, but the outgoing port does not support 802.1p tagging, the policy action will fail.
 - **DSCP/TOS** - Enable/Disable DSCP/TOS Precedence. The TOS byte is defined in RFC 791. This byte contains two fields. The precedence field is the three high-order bits (0-2) and is used to indicate the priority for the frame. The type of service field (bits 3-6) defines the throughput, delay, reliability, or cost for the frame; however, in practice these bits are not used. If you enable the **TOS Precedence** radio button, set the associated field to any value from 0-7 to specify the value that will be inserted into the precedence field of the TOS byte upon egress from the switch. A value of 7 has the highest precedence and a value of 0 has the lowest precedence. Note that you can enable **either** the DSCP or the TOS Precedence radio button to specify the mechanism you want to use (if any) to convey QoS information in the IP header of frames. DSCP and TOS are mutually exclusive. You can use either DSCP or TOS, but not both.

TCM

The TCM Policy Action option enables you to specify Three-Color Marking (TCM) actions action to impose on traffic that meets the configured policy condition(s). TCM provides a mechanism for policing network traffic by limiting the rate at which traffic is sent or received on a switch interface. TCM meters traffic based on user- configured packet rates and burst sizes and "marks" the metered packets as green, yellow, or red based on whether the traffic meets the configured rates. This "color marking" determines the packet's precedence when congestion occurs. TCM is not supported on IAP Devices and is ignored when applied to those devices.

- **Committed Traffic Policing**
 - **Committed Information Rate** - The maximum amount of bandwidth, in kbits-per-second, for all traffic that ingresses on the port.

Validity Period

The Unified Policies Validity Period Screen enables you to add a validity period to a condition by specifying the time periods when the policy is active and enforced. Four pre-configured policy validity periods are provided in the drop-down list in the **Policy Validity Periods** pane. They are **AllTheTime**, **Weekdays**, **Weekends**, and **WorkingDay**. You can also create **Custom** validity periods that are enforced during a specific timeframe.

When you have completed all of the parameters, click the **Next** button at the bottom of the screen or click on Review on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.

Note: The pre-configured validity period **AllTheTime** is the default and is automatically assigned to the condition when the **Ignore Validity Period in defining Policy Condition** checkbox is checked. You can configure a validity period when configuring and IP Condition or Service Condition. If you do not specify an IP or Service Condition, the configured period is not applied

for Wireless Controllers.

Advanced Wireless Settings

For Wireless Controllers, you can specify an absolute period or a periodic period.

- **Absolute**- Specifies an absolute time range, with a specific start and/or end time and date.
- **Periodic** - Specifies a recurring time range. Specify the start and end time and all days or selected days of the week.

Review

The Unified Policies Review Screen is used to review the Policy configuration before saving the Policy. After reviewing the Policy, click the **Create** button to save the policy to the LDAP Server. You can also click the **Back** button to return to a previous screen.

Unified Policy List

The PolicyView Unified Policy List Screen displays all configured Unified Policy Lists, including the Unified Policies included in each list, and is used to create, edit, delete, view and apply Unified Policy Lists. A Unified Policy List is a set of Unified Policies that are grouped together and assigned to devices as a group. A Unified Policy List can be applied to an AOS Switch or ClearPass Server. A Unified Policy List can be applied to wireless devices as part of an Access Role Profile. Access Role Profiles are configured in the Unified Access application (Unified Access - Device Config - Access Role Profile).

Unified Policy List Information

The following information is displayed for each Unified Policy contained in the Unified Policy List. (Click on a Unified Policy List to display the Unified Policies contained in the list.)

- **Name** - The name of the Unified Policy.
- **Condition** - The Unified Policy Condition information (e.g., IP Policy Condition would display the Source/Destination/Multicast IP address of the condition).
- **Action** - The Unified Policy Action to take if the traffic matches the Policy condition (e.g., QoS Accept/Drop)
- **Precedence** - The Precedence value of the Unified Policy (0 - 65535).
- **Validity Period** - The configured validity period for the Unified Policy.

Creating a Unified Policy List

Click on the Create icon **+** . The Create Unified Policy List Wizard appears. Complete the screens as described below, then click on the **Create** button.

Config for Policy List

Enter a Name for the Unified Policy List and select the Unified Policies you want to include in the list from the Add Unified Policies drop-down menu. (All of the currently-configured Unified Policies appear in the list. You can also click the Add icon **+** to go to Unified Policies Screen and create a new Unified Policy(ies) to add to the list.) When you select a Unified Policy from the drop-down menu, the Unified Policy will appear in a table below, so you can review the Unified Policy and modify the Precedence value, if needed.

If your are assigning a Policy List to both wired and wireless devices, select an option from the drop- down menu at the bottom of the table to override the default behavior of the devices. The default behavior for traffic that does not match a policy is different for wired and wireless devices. For wired devices, the default

behavior is to "accept" the traffic. For wireless devices, the default behavior is to "deny" the traffic. For example, if you create a Source IP Policy for a single IP address, by default wired devices would accept traffic that does not come from that IP address while wireless devices would drop the traffic. If you are assigning a Policy List to both wired and wireless devices, select an option from the drop-down menu to override device default behavior. All devices that the policy is assigned to will then follow this default behavior.

- **OV-L3-AcceptAllPolicy** - Traffic that does not match any of the policies will be accepted on all devices.
- **OV-L3-DenyAllPolicy** - Traffic that does not match any of the policies will be denied on all devices.
- **Device-Default** - Traffic that does not match any of the policies will be accepted/denied according to the device's default behavior. If you do not make a selection from the drop-down menu, this option is automatically used.

Review the Policy List configuration(s) in the table, then click the **Create** button. The new Unified Policy List will appear on the Unified Policy Lists Screen.

Note: The Wireless User Role contains a QoS rule and an Access List, which is a set of ACLs. For User Role, Wireless Controllers support two (2) QoS attributes - Bandwidth Contract - Upstream and Downstream. However, OmniVista only supports configuring Downstream Bandwidth. Additionally, the User Role can contain only a single Bandwidth Contract. So if the Unified Policy List contains more than one QoS Rule, OmniVista will display an error message: "Unified Policy List can't contain more than one QOS Action."

Device Selection

Select an option from the drop-down menu (Use Switch Picker/Use Topology), click on the **Select Devices** button and select the device(s) to which you want to apply the Policy List.

Applying a Unified Policy List to the Network


After configuring and saving a Unified Policy List, you must apply the list by notifying the devices in the network. A Unified Policy List can be assigned to AOS Switches and/or ClearPass Servers. When you click on the **Notify All** button, all of the configured Policy Lists are applied to all of the devices configured for each policy. To apply a list only to certain devices, select an option from the drop-down menu (Use Switch Picker/Use Topology), click on the **Select Devices** button and select the device(s); then click the **Notify Selected** button.


Note: A Unified Policy List can be applied to wireless devices as part of an Access Role Profile.

After notifying the devices, you can view the status of the re-cache operation, by clicking on the **Status** button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the **Audit** application, including an indication of any error that may have occurred. Note that the re-cache operation for each switch occurs in a separate thread and may take some time. Any errors that occur will also be reported in the server.txt file, in the **Audit** application.


Note: When you notify network switches, all QoS-enabled switches flush their policy tables and reload policies from the LDAP repository, which is very expensive in terms of switch resources and time. It is recommended that you review all policies that you have created and apply them at the same time to minimize switch downtime.

Editing a Unified Policy List

You can edit the Unified Policies included in a Unified Policy List or edit the Precedence value of any Unified Policy in the list. Select a Unified Policy List and click on the Edit icon . The Edit Unified Policy List Screen appears. Click on the Add Unified Policies drop-down menu. (All of the currently-configured Unified Policies appear in the list.

You can also click the Add icon  to go to Unified Policies Screen and create a new Unified policy(ies) to add to the list.) Select/unselect Unified Policies to add/remove them from the Unified Policy List. When you are finished editing the Unified Policy, click the Update button. The updated Unified Policy List will appear on the Unified Policy Lists Screen.

Deleting a Unified Policy List

To delete a Unified Policy List(s), select the list(s), click on the Delete icon , then click **OK** at the confirmation prompt. Note that you cannot delete a Unified Policy List that is associated with an Access Role Profile. To delete the list, you must first remove it from associated Access Role Profile.

Unified Policy List Information

Click on a Unified Policy List to display the information about the Policies contained in the list. The following information is displayed for each Unified Policy contained in the Unified Policy List.

- **Policy Name** - The name of the Policy.
- **Action** - The Policy Action to take if the traffic matches the Policy condition (e.g., QoS Accept/Drop)
- **Condition** - The Policy Condition information (e.g., IP Policy Condition would display the Source/Destination/Multicast IP address of the condition).
- **Precedence** - The Precedence value of the Policy (0 - 65535).
- **Validity Period** - The configured validity period for the Policy.

Resource Policies


The PolicyView Create Policy for Resources option enables you to configure Resource/Resource Group QoS Policies. Although you can use the QoS Expert Mode option to create Policies for User Network Profiles (UNP), this can be time consuming. Resources and Resource Groups can be used to quickly create Resources that can be turned into Policies and added to Policy Lists.


When you select "Create Policy From Resources" from the main PolicyView screen, the Add To Policy List Screen appears. You can use this screen to create a policy using an existing Resource/Resource Group and add it to a Policy List. If necessary, you can click on the Resource or Resource Group link on the left side of the screen to create a new Resource or Resource Group.

Add to Policy List

The PolicyView Add to Policy List Screen is used to create a policy from Resource or Resource Group and assign the policy to a Policy List.

Creating a Resource Policy

Click the checkbox next to the **Resource** or **Resource Group** Field and select an existing Resource or Resource Group from the drop-down menu (You can also click the Add icon  to go to the Resource or Resource Group screen and create a new Resource/Resource Group, before returning to this screen to select it.) Configure the remaining fields as described below, then click the **Create** button.

- **Precedence** - Set the Precedence for the Policy.
- **Action** - Set the action to be taken for the traffic if the conditions match the policy:
 - **Accept** - Accept the traffic.
 - **Drop** - Drop the traffic
 - **Platinum** - Apply Platinum Precedence to the traffic.
 - **Gold** - Apply Gold Precedence to the traffic.
 - **Silver** - Apply Silver Precedence to the traffic.
 - **Bronze** - Apply Bronze Precedence to the traffic.
- **Policy List** - Select an existing Policy List from the drop-down menu to associate with the Policy. You can also click the Add icon  to go to the Policy List screen and create a new Policy List, before returning to this screen to select it.



After creating the Resource Policy, click on the **Notify All** button to add the Resource/Resource Group to the Policy List and notify the switches in the network. To view the status of the re-cache operation, click on the **Status** button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the Audit application, including an indication of any error that may have occurred. Any errors that occur will also be reported in the server.txt file, in the Audit application. Note that the re-cache operation for each switch occurs in a separate thread and may take some time.

Important Note: Clicking the **Notify All** button causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. Any switch to which the Policy List has already been assigned will also re-cache its policy tables.

Resource


The PolicyView Resource Screen displays information for all configured Resources, and is used to create, edit, or delete a Resource. A Resource is created by specifying a name, Destination IP/Subnet, and/or a Service Group for the Resource.

Creating a Resource


Click on the Create icon . Enter a **Name** for the Resource then. Select the **Destination IP Address** checkbox and enter a **Destination IP** and **Destination Subnet**, and/or select the **Destination Service Group** checkbox and select an existing **Service Group** from the drop-down menu. (You can also click the Add icon  to go to the **Groups** application and create a new Service Group for the Resource). Click **OK** to save the Resource configuration.

Once the Resource is defined and saved, the corresponding LDAP Policy Rules, with a default initial precedence value of “50000” (which can be modified), an action of “Accept” and validity period of “AllTheTime”, is created and saved to the LDAP Server.

Editing a Resource

To edit parameters for an existing Resource, select the Resource in the Resource List and click on the Edit icon . Edit the Destination IP Address and/or Destination Service Group parameter(s), then click **OK**. You cannot edit the Resource Name, you can only delete the Resource and create a new one.



Deleting a Resource

To delete a Resource, select the Resource(s) in the Resource List and click on the Delete icon , then click **OK** at the confirmation prompt.

Resource Group


The PolicyView Resource Group Screen displays information for all configured Resource Groups, and is used to create, edit, or delete a Resource Group from existing Resources.

Creating a Resource Group

Click on the Create icon . Enter a Name for the Resource Group and select a Resource(s) from the list of Resources. (You can also click the Add icon  to go to the Resource Screen and create a new Resource, before returning to this screen to select it.) Click OK to save the Resource Group configuration.

Once the Resource Group is defined and saved, the corresponding LDAP Policy Rules, with a default initial precedence value of “50000” (which can be modified), an action of “Accept” and validity period of “AllTheTime”, is created and saved to the LDAP Server.

Editing a Resource Group

To edit parameters for an existing Resource Group, select the Resource Group in the Resource Group List and click the Edit icon . Add/delete Resources to/from the Resource group, then click OK. You cannot edit the Resource Group Name, you can only delete the Resource Group and create a new one.

Deleting a Resource Group

To delete a Resource Group, select the Resource Group(s) in the Resource Group List and click on the Delete icon , then click **OK** at the confirmation prompt.

One Touch Policies

The PolicyView QoS One Touch option enables you to quickly create One Touch Data, ACL, and Voice Policies for network traffic. One Touch Data Policies enable you to assign a desired quality of service - Platinum, Gold, Silver, or Bronze - to all traffic flowing to, and originating from, specific Data Servers. One Touch ACL Policies enable you to create ACL Policies for all traffic flowing to, and originating from, specific Network Groups. All policies created are applied to all QoS-enabled devices in the List of All Discovered Devices. (You can use the Expert mode if you need to assign a different priority to any server.)

When you select "One Touch Policy" from the main PolicyView screen, the One Touch Data Screen appears. Click on the One Touch ACLs link on the left side of the screen to configure One Touch ACL Policies. Click the One Touch Voice link on the left side of the screen to configure One Touch Voice Policies.

One Touch Data Policies



The PolicyView One Touch Data Screen displays all of the Data Servers that have been configured with a One Touch Data Policy and the status of the policy on the LDAP Server:

- **Saved** - The policy has been successfully written to the LDAP Server.
- **Unsaved** - The new policy, or modified policy has not been saved to the LDAP Server
- **Error** - An error condition has made it impossible to write the policy to the LDAP Server.

The screen is used to create, edit, or delete one of the following One Touch Data Policies:

- **Platinum** - provides the highest quality of service (and maps to a firmware priority of 7)
- **Gold** - provides the next-highest quality of service (and maps to a firmware priority of 5)
- **Silver** - provides the next-highest quality of service (and maps to a firmware priority of 3)
- **Bronze** - provides the same quality of service as best effort (and maps to a firmware priority of 1)

Creating a One Touch Data Policy


Select a Priority from the drop-down menu, then click on the Create icon . Enter the IP address of the Data Server in the **Server IP Address** field, and click **Create**. The Data Server will appear in the table with the Status "Unsaved". Click on the Save icon  to save the Policy to the LDAP Server. The Priority selected will be applied to all Data Servers in the table. When you are finished, click on **Notify All** to apply the policy to all of the switches in the network.

Important Note: Clicking the **Notify All** button causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. If any One Touch Data policy has already been defined, the switch(es) to which the policy is assigned will also re-cache its policy tables. **It is recommended that you verify all policies that you have created and apply them at the same time to minimize switch downtime.**

Applying a One Touch Data Policy


When you click the **Notify All** button, the policy(ies) is applied to all switches in the network. To view the status of the re-cache operation, click on the **Status** button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the **Audit** application, including an indication of any error that may have occurred. Any errors that occur will also be reported in the server.txt file, in the **Audit** application. Note that the re-cache operation for each switch occurs in a separate thread and may take some time.

Editing a One Touch Data Policy

Select a Priority from the drop-down menu. The status of all of the Data Servers in the table will change to "Unsaved". Click on the Save icon  to save the Policy to the LDAP Server. The Priority selected will be applied to all Data Servers in the table. When you are finished, click on **Notify All** to apply the policy to all of the switches in the network.

Important Note: Clicking the **Notify All** button causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. If any One Touch Data policy has already been defined, the switch(es) to which the policy is assigned will also re-cache its policy tables. **It is recommended that you verify all policies that you have edited and apply them at the same time to minimize switch downtime.**

Deleting One Touch Data Policy

To delete a One Touch Data Policy(ies), select the server(s) in the table and click on the Delete icon , then click **OK** at the confirmation prompt. When you click the **OK** button:

- All One Touch Data policies for the servers you selected are removed from the LDAP Server.
- All One Touch Data policies for the servers you selected are removed from switch attributes in the LDAP "role" objects.

- The server(s) are removed from the table.
- A confirmation message is displayed when the LDAP Server has been successfully updated. The success of the operation is also reported in the policy.log file in the **Audit** application.
- An SNMP message is sent to each QoS-qualified device in the List of All Discovered Devices, informing that the information in the LDAP repository has changed and commanding the devices to update their cached policies with the current information from the LDAP repository. Note that the re-cache operation for each switch occurs in a separate thread and may take some time. Be sure to check the policy.log file in the Audit application for the re-cache status of each switch.

Example of a New One Touch Policy (or Edit Policy) Operation

Let's say **Platinum** was selected as the priority for the following two Server IP addresses:

164.178.32.107

164.178.33.51

When saved, the following policies are created and written to the LDAP Server:

OneTouchDR\$\$164.178.32.107

Condition specifies traffic originating from source IP address 164.178.32.107

Action specifies Platinum QoS for this traffic

OneTouchDR\$D164.178.32.107

Condition specifies traffic transmitted to destination IP address 164.178.32.107

Action specifies Platinum QoS for this traffic

OneTouchDR\$\$164.178.33.51

Condition specifies traffic originating from source IP address 164.178.33.51

Action specifies Platinum QoS for this traffic

OneTouchDR\$D164.178.33.51

Condition specifies traffic transmitted to destination IP address 164.178.33.51

Action specifies Platinum QoS for this traffic

Please note that the names beginning with "OneTouchDR" are the names used for the policies in the LDAP Server. Within the PolicyView QoS application, all One Touch Data policies are referred to by the generic composite name **OneTouchDR**, no matter how many individual One Touch Data policies have been written to the LDAP Server. One Touch Data rules that have been created automatically by PolicyView can be viewed in the Expert mode window.



One Touch ACL Policies

The PolicyView One Touch ACLs Screen displays all of the Network Groups that have been configured with a One Touch ACL Policy, as well as the Accessibility configured for the policy (Accept/Drop), and status of the policy on the LDAP Server.


- **Saved** - The policy has been successfully written to the LDAP Server.
- **Unsaved** - The new policy, or modified policy has not been saved to the LDAP Server
- **Error** - An error condition has made it impossible to write the policy to the LDAP Server.

The screen is used to create, edit, or delete a One Touch ACL Policy for a Network Group.

Creating a One Touch ACL Policy

Click on the Create icon . Select an existing Network Group from the **ACL IP Server Group** drop-down menu, and an **Accessibility** option for the group (Accept/Drop), and click **Create**. The One Touch ACL Policy will appear in the table with the Status "Unsaved". Click on the Save icon  to save the Policy to the LDAP Server. The Priority selected will be applied to all Data Servers in the list. When you are finished, click on the **Notify All** button to apply the policy to all of the switches in the network.



Important Note: Clicking the **Notify All** button causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. If any One Touch ACL Policy has already been defined, the switch(es) to which the policy is assigned will also re-cache its policy tables. **It is recommended that you verify all policies that you have created and apply them at the same time to minimize switch downtime.**

Note: You can also click the Add icon  to go to the **Groups** application and create a new Network Group, before returning to this screen to select it.

Applying a One Touch ACL Policy


When you click the **Notify All** button, the policy(ies) is applied to all switches in the network. To view the status of the re-cache operation, click on the **Status** button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the **Audit** application, including an indication of any error that may have occurred. Any errors that occur will also be reported in the server.txt file, in the **Audit** application. Note that the re-cache operation for each switch occurs in a separate thread and may take some time.

Editing a One Touch ACL Policy

Select the policy and click the Edit icon . Edit the **Accessibility** Field (Accept/Drop), click **Update**. The policy will appear in the table with the Status "Unsaved". If necessary, repeat to edit additional entries in the list. You cannot edit the ACL IP Server Group. When you are finished, click on the Save icon  to save the update(s) to the LDAP Server, then click on the **Notify All** button to apply the policy to all of the switches in the network.

Important Note: Clicking **Notify All** causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. If any One Touch ACL policy has already been defined, the switch(es) to which the policy is assigned will also re-cache its policy tables. **It is recommended that you verify all policies that you have edited and apply them at the same time to minimize switch downtime.**

Deleting a One Touch ACL Policy

To delete a policy(ies), select the policy(ies) in the table and click on the Delete icon , then click **OK** at the confirmation prompt. When you click the **OK** button:

- All One Touch ACL policies for the ACL IP Network Group(s) you selected are removed from the LDAP Server.
- All One Touch ACL policies for the ACL IP Network Group(s) you selected are removed from switch attributes in the LDAP "role" objects.
- The Network Groups are removed from the table.
- A confirmation message is displayed when the LDAP Server has been successfully updated. The success of the operation is also reported in the policy.log file in the Audit application.

- An SNMP message is sent to each QoS-qualified device in the List of All Discovered Devices, informing them that the information in the LDAP Server has changed and commanding the devices to update their cached policies with the current information from the LDAP Server. Note that the re-cache operation for each switch occurs in a separate thread and may take some time. Be sure to check the policy.log file in the Audit application for the re-cache status of each switch.

Example of a One Touch ACL Policy Creation

Let's say we have selected **Data Center Switches** as the Network Group and **Accept** as the accessibility option. When saved, the following policies are created and written to the LDAP Server:

OneTouchAR\$\$Data Center Switches

Condition specifies traffic originating from source IP Network group 'data center switches'
Action specifies accept as the disposition for this traffic

OneTouchAR\$\$Data Center Switches

Condition specifies traffic transmitted to destination IP Network group 'data center switches'
Action specifies accept as the disposition for this traffic

Note: Names beginning with "OneTouchAR" are the names used for the One Touch ACL policies in the LDAP repository.

One Touch Voice Policies

The PolicyView One Touch Voice Policies enable you to easily assign the highest quality of service to all voice traffic that is destined for Alcatel-Lucent Enterprise devices. There are four QoS priority queues supported by Alcatel-Lucent Enterprise devices: Platinum, Gold, Silver, and Bronze. Platinum provides the highest QoS and Bronze provides the lowest QoS. One Touch Voice Policies enable you to assign **Platinum** QoS to voice traffic. You can assign the policies by IP Subnet or MAC Address:

- One Touch Voice IP Policies - Create Platinum Layer 3 Policies for all voice traffic flowing to, and originating from, an IP subnet.
- One Touch Voice MAC Policies - Create Platinum Layer 2 Policies for all voice traffic flowing to, and originating from, devices for a specific vendor by entering a range of MAC Addresses for the vendor.

One Touch Voice IP Policies

The PolicyView One Touch Voice IP Policies Screen displays all of the One Touch Voice IP Policies that have been configured and the status of the policy on the LDAP Server:

- **Saved** - The policy has been successfully written to the LDAP Server.
- **Unsaved** - The new policy, or modified policy has not been saved to the LDAP Server
- **Error** - An error condition has made it impossible to write the policy to the LDAP Server.

The screen is used to create, edit, or delete One Touch Voice IP Policies for all voice traffic flowing to, and originating from, an IP subnet.

Creating a One Touch Voice IP Policy

Click on the Create icon **+**. Enter a **Subnet IP** and **Subnet Mask** and click **Create**. The Network Subnet will appear in the table with the Status "Unsaved". If necessary, click on the Create icon **+** to create additional policies. When you are finished, click on the Save icon **✓** to save the Policy(ies) to the LDAP



Server. Next, click on the **Notify All** button to apply the policy to all of the switches in the network. Note that when you create a One Touch Voice IP Policy, two policies are created for each subnet entered - one for traffic originating from the subnet, and one for traffic flowing to the subnet.

Important Note: Clicking **Notify All** causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. If any One Touch Voice IP Policy has already been defined, the switch(es) to which the policy is assigned will also re-cache its policy tables. **It is recommended that you verify all policies that you have created and apply them at the same time to minimize switch downtime.**

Applying a One Touch Voice IP Policy

When you click the **Notify All** button, the policy(ies) is applied to all switches in the network. To view the status of the re-cache operation, click on the **Status** button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the Audit application, including an indication of any error that may have occurred. Any errors that occur will also be reported in the server.txt file, in the Audit application. Note that the re-cache operation for each switch occurs in a separate thread and may take some time.


Editing a One Touch Voice IP Policy

Select a Subnet from the One Touch Voice IPs List and click the Edit icon . Edit the **Subnet Mask**, click **Update**, then click on the Save icon  to save the update to the LDAP Server. If necessary, repeat to edit additional entries in the list. You cannot edit the IP Subnet you can only delete it and create a new one.

When you are finished, click on the **Notify All** button to apply the policy to all of the switches in the network.

Important Note: Clicking **Notify All** causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. If any One Touch Voice IP Policy has already been defined, the switch(es) to which the policy is assigned will also re-cache its policy tables. **It is recommended that you verify all policies that you have created and apply them at the same time to minimize switch downtime.**

Deleting a One Touch Voice IP Policy

To delete a policy(ies), select the Subnet IP in the list, click on the Delete icon , then click **OK** at the confirmation prompt. When you click the **OK** button, all One Touch Voice IP Policies for the IP Subnet(s) you deleted are removed from the LDAP Server (both Layer 2 and Layer 3 Policies). When you are finished, click on the **Notify All** button to apply the deletion to all of the switches in the network. .

Important Note: Clicking **Notify All** causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. If any One Touch Voice IP Policy has already been defined, the switch(es) to which the policy is assigned will also re-cache its policy tables. **It is recommended that you verify all policies that you have created and apply them at the same time to minimize switch downtime.**




One Touch Voice MAC Policies

The PolicyView One Touch Voice MAC Policies Screen displays all of the One Touch Voice MAC Policies that have been configured and the status of the policy on the LDAP Server:

- **Saved** - The policy has been successfully written to the LDAP Server.
- **Unsaved** - The new policy, or modified policy has not been saved to the LDAP Server
- **Error** - An error condition has made it impossible to write the policy to the LDAP Server.

The screen is used to create, edit, or delete One Touch Voice MAC Policies for all voice traffic flowing to, and originating from, a specific vendor by entering a range of MAC Addresses for the vendor.

Creating a One Touch Voice MAC Policy

Click on the Create icon . Enter a **Vendor** Name and **MAC Prefix** and click **Create**. The policy will appear in the One Touch Voice MACs List with the Status "Unsaved". If necessary, click on the Create icon  to create additional policies. When you are finished, click on the Save icon  to save the Policy(ies) to the LDAP Server. Next, click on the **Notify All** button to apply the policy to all of the switches in the network.



Note that when you create a One Touch Voice MAC Policy, two policies are created for each MAC Address entered - one for traffic originating from the MAC Address, and one for traffic flowing to the MAC Address.

Important Note: Clicking **Notify All** causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. If any One Touch Voice MAC Policy has already been defined, the switch(es) to which the policy is assigned will also re-cache its policy tables. **It is recommended that you verify all policies that you have created and apply them at the same time to minimize switch downtime.**

Applying a One Touch Voice MAC Policy


When you click the **Notify All** button, the policy(ies) is applied to all switches in the network. To view the status of the re-cache operation, click on the **Status** button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the Audit application, including an indication of any error that may have occurred. Any errors that occur will also be reported in the server.txt file, in the **Audit** application. Note that the re-cache operation for each switch occurs in a separate thread and may take some time.

Editing a One Touch Voice MAC Policy

Select a Policy from the One Touch Voice MACs List and click the Edit icon . Edit the **MAC Prefix**, click **Update**, then click on the Save icon  to save the update to the LDAP Server. If necessary, repeat to edit additional entries in the list. You cannot edit the **Vendor** Name. You can only delete it and create a new one. When you are finished, click on the **Notify All** button to apply the policy to all of the switches in the network.

Important Note: Clicking **Notify All** causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. If any One Touch Voice MAC Policy has already been defined, the switch(es) to which the policy is assigned will also re-cache its policy tables. **It is recommended that you verify all policies that you have created and apply them at the same time to minimize switch downtime.**

Deleting a One Touch Voice MAC Policy

To delete a policy(ies), select the Policy in the list, click on the Delete icon , then click **OK** at the confirmation prompt. When you click the **OK** button, all One Touch Voice MAC Policies you deleted are removed from the LDAP Server (both Layer 2 and Layer 3 Policies). When you are finished, click on the

Notify All button to apply the deletion to all of the switches in the network.

Important Note: Clicking **Notify All** causes all QoS-enabled switches to flush their policy tables and reload policies from the LDAP Server, which is very expensive in terms of switch resources and time. If any One Touch Voice MAC Policy has already been defined, the switch(es) to which the policy is assigned will also re-cache its policy tables. **It is recommended that you verify all policies that you have created and apply them at the same time to minimize switch downtime.**

View/Modify Policies and Policy List

The PolicyView View/Modify Policies and Policy List option enables you to view existing Policies, view/create/modify Policy Lists, and view Policies for specific switches.

- The Policies Screen displays all configured Policies. You can also click on a policy to view detailed information about the Policy (e.g., Condition, Action).
- The Policy Lists Screen displays all configured Policy Lists, along with their Policies, and is used to add/edit/delete Policy Lists.
- The Policies by Switch Screen enables you to view Policies for specific switches in the network.

Policies

The PolicyView Policies Screen displays basic information for all configured Policies. You can also click on a policy to view detailed information about the Policy (e.g., Condition, Action).

Policy Information

The list below defines all of the possible fields that can be displayed in the table. Not all of the fields below are displayed in the Default view. However, you can click on the **Select All** button to display all fields; or click on the **Custom** button to create a custom view to display only specified fields. (Click on the **Custom** button, then click on the **Custom Template** button, select the fields you want to display, then click **OK**.) After creating a custom view, you can just click on the **Custom** button to display that custom view. You can change the custom view at any time for a different display.

- **Policy Name** - The name of the Policy.
- **Scope** - The scope of the Policy (e.g., IP Filtering, Provisioned).
- **Status** - Indicates whether or not the Policy has been saved to the LDAP Server (Saved/Unsaved).
- **Precedence** - The Precedence value of the Policy (0 - 65535).
- **Enabled** - Indicates whether or not the Policy is enabled (Yes/No).
- **Save** - Indicates whether or not the rule will be recorded during a snapshot command (Yes/No).
- **Log Matches** - Indicates whether or not matches to this rule are logged in the QoS Log (Yes/No).
- **Reflexive** - Indicates whether or not the Policy is reflexive. Reflexive Policy Rules allow specific return connections that would normally be denied (Yes/No).
- **Default List** - Indicates whether or not the Policy is saved to the Default List. A Default Policy List always exists in the switch configuration. By default, a Policy Rule is added to this list when the rule is created. A rule remains a member of the Default List even when it is subsequently assigned to additional lists (Yes/No).
- **SLA Policy Trap** - Indicates whether or not an SLA Policy Trap is configured for the policy.



Detailed Policy Information

- **Policy Rule** - The name of the Policy Rule and the Policy Rules configured for the Policy.
- **Policy Condition** - The Policy condition information (e.g., IP Policy condition would display the Source/Destination/Multicast IP address of the condition).
- **Policy Action** - The Policy action to take if the traffic matches the Policy condition (e.g., QoS Accept/Drop)
- **Policy Validity Period** - The configured validity period for the Policy.
- **Policy Roles** - The switches to which the Policy has been assigned.



Policy Lists

The PolicyView Policy Lists Screen displays all configured Policy Lists, including the Policies Rules included in each list, and is used to create, edit, and delete Policy Lists. You can click on a Policy List to display the Policies contained in the list. A Policy List is a set of Policies that are grouped together and can be assigned to switches as a group. The QoS/ACL Policies that you add to a Policy List can be defined using the PolicyView Application. You can also include Resources or Resource Groups in a Policy using the "Resources - Add to Policy" Screen.

Creating a Policy List


Click on the Create icon . The Create Policy List Screen appears. Enter a **Name** for the Policy List and select the Policies you want to include in the list from the **Add Policies** drop-down menu. (All of the currently- configured Policies appear in the list. You can also click the Add icon  to go to Expert Mode and create a new policy(ies) to add to the list.) When you select a Policy from the drop-down menu, the Policy will appear in a table below, so you can review the Policy and modify the Precedence value, if needed. When you are finished reviewing the Policy(ies), click the **Create** button. The new Policy List will appear on the Policy Lists Screen.

Editing a Policy List

You can edit the Policies included in a Policy List or edit the Precedence value of any Policy in the list. Select a Policy List and click on the Edit icon . The Edit Policy List Screen appears. Click on the **Add Policies** drop-down menu . (All of the currently-configured Policies appear in the list. You can also click the Add icon  to go to Expert Mode and create a new policy(ies) to add to the list.) Select/unselect Policies to add/remove them from the Policy List. When you are finished editing the Policy, click the **Update** button. The updated Policy List will appear on the Policy Lists Screen.

Note: Adding/deleting a policy to/from a policy list will automatically update any switch roles that contain this policy list with the updated policies.

Deleting a Policy List

To delete a Policy List(s), select the list(s), click on the Delete icon , then click **OK** at the confirmation prompt. Note that you cannot delete a Policy List that is associated with a User Network Profile (UNP). To delete the list, you must first remove it from the UNP.

Policy List Information

Click on a Policy List to display the Policies contained in the list. The following information is displayed for each Policy contained in the Policy List.

- **Name** - The name of the Policy.
- **Condition** - The Policy Condition information (e.g., IP Policy Condition would display the Source/Destination/Multicast IP address of the condition).
- **Action** - The Policy Action to take if the traffic matches the Policy condition (e.g., QoS Accept/Drop)
- **Precedence** - The Precedence value of the Policy (0 - 65535).
- **Validity Period** - The configured validity period for the Policy.

Policies by Switch

The PolicyView Policies by Switch Screen enables you to view Policies configured on specific devices. To view Policies for specific devices, click on the **Select Device** button to bring up the Selection Wizard. Select the device(s) you want to display, then click **OK**. Device information and Policy Status for each device is displayed. Select a device to display Policy and Policy List information for the device.

Device Information

The following information is displayed for each selected device:

- **Name** - The user-configured name for the device.
- **Address** - The IP Address of the device.
- **DNS Name** - The DNS name of the device, if applicable.
- **Device Type** - The device model (e.g., OS6850-48, OS6900-X20).
- **Version** - The AOS software version running on the device (e.g., 732.344.R01).
- **Policy Status** - Indicates whether or not the switch has re-cached its Policy information to contain the Policy. After a Policy is created and saved on the LDAP Server, the user assigns the Policy to the switch. This causes the switch to flush its Policy Tables and reload the latest Policies from the LDAP Server.

Policy Information

The following Policy information is displayed when you click on a device (click on a Policy to display detailed information about the Policy):

- **Name** - The name of the Policy.
- **Scope** - The scope of the Policy (e.g., IP Filtering, Provisioned).
- **Status** - Indicates whether or not the Policy has been saved to the LDAP Server (Saved/Unsaved).
- **Precedence** - The Precedence value of the Policy (0 - 65535).
- **Enabled** - Indicates whether or not the Policy is enabled (Yes/No).
- **Save** - Indicates whether or not the rule will be recorded during a snapshot command (Yes/No).
- **Log Matches** - Indicates whether or not matches to this rule are logged in the QoS Log (Yes/No).
- **Reflexive** - Indicates whether or not the Policy is reflexive. Reflexive Policy Rules allow specific return connections that would normally be denied (Yes/No).
- **Default List** - Indicates whether or not the Policy is saved to the Default List. A Default Policy List always exists in the switch configuration. By default, a Policy Rule is added to this list when the rule is created. A rule remains a member of the Default List even when it is subsequently assigned to additional lists (Yes/No).
- **SLA Policy Trap** - Indicates whether or not an SLA Policy Trap is configured for the policy.

Detailed Policy Information

The following information is displayed when you click on a Policy:

- **Policy Rule** - The name of the Policy Rule configured for the Policy.
- **Policy Condition** - The Policy condition information (e.g., IP Policy condition would display the Source/Destination/Multicast IP address of the condition).
- **Policy Action** - The Policy action to take if the traffic matches the Policy condition (e.g., QoS Accept/Drop)
- **Policy Validity Period** - The configured validity period for the Policy.
- **Policy Roles** - The switches to which the Policy has been assigned.

Policy List Information

The following Policy List information is displayed when you click on a device:


- **Name** - The Policy List name.
- **Policies** - The Policy Rules configured for the Policy.

Expert Mode

The PolicyView Expert Mode option is used to create, edit, delete, and view custom QoS Policies. Custom Policies are created using a wizard that guides you through each of the steps needed to create the Policy and apply the Policy to switches in the network. All currently-configured Policies are listed in the Existing Policies Table. You can view basic information about a Policy in the table or click on a specific policy to view more detailed information.

Note: You cannot create, delete, or edit a One Touch Policy in the Expert mode. You must use the applicable One Touch option (Data, ACL, Voice) to create, delete, or edit a One Touch policy.

Creating a Custom Policy

Custom Policies are created using a wizard that guides you through each of the steps needed to create the policy and apply the policy to switches in the network. To create a Custom Policy, click on the Create icon . The wizard will then guide you through the following screens:

- Configuration - Basic policy configuration (e.g., Policy Name, Precedence)
- Device Selection - Specify the devices to which you will apply the policy
- Set Condition - Specify the conditions that must be true before traffic will be allowed to flow.
- Set Action - Specify parameters for the traffic that will flow.
- Validity Period - Specify the time period for the policy to be in effect.
- Review - Review the policy details before creating the policy.

Note: As you configure a policy, conditions and actions are verified against the devices selected for the policy. If a condition or action is not supported by one of the selected devices, an error message will appear indicating the error and corrective action to be taken.


Applying a Custom Policy to the Network

After configuring and saving a policy(ies), you must apply the policy(ies) by notifying the switches in the network. When you click on the **Notify All** button, all of the policies listed in the Existing Policies Table are applied to all of the switches configured for each policy. To apply the policy(ies) only to certain devices, select an option from the drop-down menu (Use Switch Picker/Use Topology), click on the **Select Device** button and select the device(s); then click the **Notify Selected** button.

After notifying the switches, you can view the status of the re-cache operation, by clicking on the **Status** button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the Audit application, including an indication of any error that may have occurred. Note that the re-cache operation for each switch occurs in a separate thread and may take some time. Any errors that occur will also be reported in the server.txt file, in the Audit application.

Note: When you notify network switches, all QoS-enabled switches flush their policy tables and reload policies from the LDAP repository, which is very expensive in terms of switch resources and time. It is recommended that you review all policies that you have created and apply them at the same time to minimize switch downtime.

Editing a Custom Policy

To edit a policy, select the policy in the Existing Policy Table and click on the Edit icon . Use the wizard to make any edits. When you are done, apply the edited policy to the network.

Deleting a Custom Policy

To delete a policy(ies), select the policy(ies) in the Existing Policies Table and click on the Delete icon , then click **OK** at the confirmation prompt.

Policy Information

The Existing Policies Table displays basic information for all configured Policies. You can also click on a policy to view detailed information about the Policy (e.g., Condition, Action). Note that Unified Policies are not displayed in the Expert Mode Existing Policies Table. They are only displayed in the Existing Unified Policies Table.

Basic Information

- **Policy Name** - The name of the Policy.
- **Scope** - The scope of the Policy (e.g., IP Filtering, Provisioned).
- **Precedence** - The Precedence value of the Policy (0 - 65535).
- **Status** - Indicates whether or not the Policy has been saved to the LDAP Server.
- **Enable** - Indicates whether or not the Policy is enabled.
- **Save** - Indicates whether or not the rule will be recorded during a snapshot command.
- **Log Matches** - Indicates whether or not matches to this rule are logged in the QoS Log.
- **Reflexive** - Indicates whether or not the Policy is reflexive. Reflexive Policy Rules allow specific return connections that would normally be denied (Yes/No).
- **Default List**- Indicates whether or not the Policy is saved to the Default Policy List. By default, a Policy Rule is added to this list when it is created. A Policy Rule remains a member of the Default List even when it is subsequently assigned to additional Policy Lists.
- **SLA Policy Trap** - Indicates whether or not an SLA Policy Trap is configured for the policy.

Detailed Information

- **Policy Rule** - The name of the Policy Rule and the Policy Rules configured for the Policy.
- **Policy Condition** - The Policy condition information (e.g., IP Policy condition would display the Source/Destination/Multicast IP address of the condition).
- **Policy Action** - The Policy action to take if the traffic matches the Policy condition (e.g., QoS Accept/Drop)
- **Policy Validity Period** - The configured validity period for the Policy.
- **Policy Roles** - The switches to which the Policy has been assigned.

Config for Policy

The Expert Mode Config for Policy Screen is used to configure basic Policy parameters. When you have completed all of the parameters, click the **Next** button at the bottom of the screen or click on Device Selection on the left side of the screen to move to the next step.

- **Name** - The Policy name.
- **Precedence** - The Policy precedence. By default, the precedence field is pre-filled with the lowest unused precedence value (Range = 0 - 65535).

Click on **Show Advanced Options** to display and configure the options below. By default, these options are set to **Ignore**.

- **Default List** - Adds the rule to the QoS Default Policy List.
- **Enable** - Enables the policy.
- **Save** - Marks the policy rule so that it may be captured as part of the switch configuration.
- **Log Matches** - Configures the switch to log messages about specific flows coming into the switch that match this policy rule.
- **Send Trap** - Enables traps for the Policy.
- **Reflexive** - Enables support for the Reflexive for the policy. Reflexive policies allow specific return connections that would normally be denied.

Device Selection

The Expert Mode Device Selection Screen is used to select the switches to which you want to apply the Policy. Select an option (Use Switch Picker/Use Topology), and select the device(s). Click on the **Next** button at the bottom of the screen or click on Set Condition on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen and add/delete devices.

Set Condition

The Expert Mode Set Condition Screen contains a list of Conditions that you can configure for the Policy (e.g., Interface Condition, MAC Condition). When you create a Condition, the Condition(s) you configure must be true before traffic is allowed to flow. Click on a Condition to display the configuration options for the Condition. (Click again on the Condition to close the configuration options.) When you have completed all of the parameters for the Condition(s), click the **Next** button at the bottom of the screen or click on Set Action on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.

Conditions

A brief description of each Condition is provided below. Click the hyperlink for each Condition for detailed configuration instructions.

- **L1 Interfaces** - Create a Condition that applies the policy to traffic flowing from a specific source interface type or to traffic flowing to a specific destination interface type.
- **L2 MACs** - Create a Condition that applies the policy to traffic originating from a MAC address/group or to traffic flowing to a MAC address/group. (Note that any MAC address may contain wildcard characters).
- **L2 VLANs** - Create a Condition that applies the policy to traffic flowing from a source VLAN to a destination VLAN, or to traffic flowing from one source VLAN to any destination VLAN, or to traffic flowing from any source VLAN to one destination VLAN.
- **L2 802.1P** - Create a Condition that applies the policy to traffic with a specified 802.1 priority value.
- **L3 IPs** - Create a Condition that applies the policy to traffic originating from an IP address/network group or to traffic flowing to an IP address/network group. (Note that any IP address can be masked).
- **L3 DSCP/TOS** - Create a Condition that applies the policy to traffic with a specified value in either the DSCP (Differentiated Services Code Point) byte or in the IP TOS (IP Type of Service) byte. Both DSCP and IP TOS are mechanisms used to convey QoS information in the IP header of frames.
- **L3 TCP Flags** - Creates a Condition that applies the policy to traffic based on TCP values.
- **L4 Services** - Create a Condition that applies the policy to traffic flowing between two TCP or UDP ports, or to all traffic originating from a TCP or UDP port, or to all traffic flowing to a TCP or UDP port. You can also create a Condition using an existing service/service group.
- **L7 Applications** - Create a Condition that applies the policy to traffic associated with a specific Application Group.
- **Application Visibility** - Create a condition that applies the policy to traffic associated with a specific Application Group. Application Name Conditions are **not** supported at this time.
- **VXLAN** - Create a VM Snooping Condition that applies to incoming VXLAN packets.

Note: Please refer to the switch Release Notes for information on the specific QoS functions available on various platforms and combinations of hardware/firmware.

L1 Interfaces


An Interface Condition applies the Policy to a traffic flowing from/to an interface type. Select the parameter(s) you want to configure by selecting the applicable checkbox, then select an option from the drop-down menu.

- **Source Interface** - Selecting a Source Interface type, restricts the policy to a traffic type that flows from that interface type only. If you do not select this option, you are effectively stating that the source traffic type is not a criterion for the Policy.
- **Destination Interface** - Selecting a Destination Interface, restricts the policy to a traffic type that flows to that interface type only. If you do not select this option, you are effectively stating that the destination traffic type is not a criterion for the policy.
- **Other Type** - Entering an Ethernet Type, restricts the policy to this type of ethernet traffic. If you do not select this option, you are effectively stating that the ethernet type is not a criterion for the policy.

L2 MACs

A MAC Condition applies the Policy to traffic flowing from/to a MAC Address/Group. Note that Layer 2 Conditions (conditions that specify MAC Addresses) are "lost" when traffic passes through a router. For this reason, it may be advisable to specify other types of Conditions (such as a Layer 3 Condition, which

specifies IP Addresses) when traffic is expected to travel more than one router hop.

Select the parameter(s) you want to configure by selecting the applicable checkbox. Click on **Single** to configure a single MAC Address or **Group** to configure a MAC Group, then enter a MAC address or select a MAC Group from the drop-down menu. (You can also click the Add icon  to go to the **Groups** application and create a new MAC Group.)


- **Source MAC Address/MAC Group** - Configuring a Source MAC Address/Group Condition restricts the policy to traffic that flows from this MAC Address/Group only. If you do not select this option, you are effectively stating that the Source MAC Address/Group traffic is not a criterion for the policy.
- **Destination MAC Address/MAC Group** - Configuring a Destination MAC Address/Group Condition restricts the policy to traffic that flows to this MAC Address/Group only. If you do not select this option, you are effectively stating that the Destination MAC Address/Group traffic is not a criterion for the policy.

Notes:

- Conditions that specify both a source and a destination MAC address may be rejected by some switch platforms as invalid. However, if you wish to create policies for both source and destination traffic, you can create one policy for the source traffic and a second policy for the destination traffic.
- MAC addresses may contain the wildcard character *. However, one * character must be entered for each individual hex digit in the MAC address: for example, **00435C:*******, not **00435C:***.
- The following MAC address ranges are assigned to Alcatel-Lucent Enterprise voice devices and Alcatel-Lucent Enterprise IP phones. You can create Conditions specifying these address ranges using the MAC Address tab.
 - Voice Devices
 - 00809F3A0000 - 00809F3AFFFF
 - 00809F3B0000 - 00809F3BFFFF
 - 00809F3C0000 - 00809F3CFFFF
 - IP Phones
 - 00809F3D0000 - 00809F3DFFFF
 - Multi-Media Devices
 - 00809F3E0000 - 00809F3EFFFF
 - 00809F3F0000 - 00809F3FFFFF

L2 VLANs

A VLAN Condition applies the Policy to traffic flowing from/to a VLAN/VLAN Group. You can also create an Inner Source VLAN Condition for a stacked VLAN network, and a Condition based on Virtual Routing and Forwarding (VRF) name (OS10K).

Select the parameter(s) you want to configure by selecting the applicable checkbox. For VLANs/VLAN Groups, click on **Single** to configure a single VLAN or **Group** to configure a VLAN Group, then enter a VLAN or select a VLAN Group from the drop-down menu. (You can also click the Add icon  to go to the **Groups** application and create a new VLAN Group.)

- **Source VLAN/VLAN Group** - Configuring a Source VLAN/VLAN Group Condition restricts the policy to traffic that flows from this VLAN/VLAN Group only. If you do not select this option, you are effectively stating that the Source VLAN/VLAN Group traffic is not a criterion for the policy.
- **Destination VLAN/VLAN Group** - Configuring a Destination VLAN/VLAN Group Condition restricts the policy to traffic that flows to this VLAN/VLAN Group only. If you do not select this option, you are

effectively stating that the Destination VLAN/VLAN Group traffic is not a criterion for the policy.

- **Inner Source VLAN** - An Inner Source VLAN Condition is applied to double-tagged VLAN Stacking traffic and is used to classify such traffic based on the inner VLAN ID tag, also known as the customer VLAN ID. Configuring an Inner Source VLAN Condition restricts the policy rule to all double-tagged traffic for that VLAN. If you do not select this option, you are effectively stating that the Inner Source VLAN traffic is not a criterion for the policy.
- **VRF Name** - Configuring a VRF Name Condition restricts the policy to traffic that flows to this VRF only. If you do not select this option, you are effectively stating that VRF traffic is not a criterion for the policy. Note that by default, QoS Policy Conditions are not associated with any specific VRF instance. The Policy applies across all instances.

L2 802.1P


An 802.1P Condition applies the Policy to traffic that has a specified 802.1 priority value in the header of the frame. 802.1p is the IEEE extension of 802.1d and is a standard for the use of MAC-layer bridges in filtering and expediting multicast traffic. 802.1p prioritizes traffic through the insertion of a three-bit priority value into the header of the frame. An 802.1 priority value of 7 provides the highest priority, and an 802.1 priority value of 0 provides the lowest priority. Select the parameter(s) you want to configure by selecting the applicable checkbox, then enter a priority value.

- **802.1 Priority Level** - Set the field to the desired priority value (0-7). This will restrict the policy to incoming traffic that has that 802.1 Priority value in the frame header. A value of 7 provides the highest priority and a value of 0 provides the lowest priority. If you do not select this option, you are effectively stating that the 802.1P Priority Level is not a criterion for the Policy.
- **Inner 802.1 Priority Level** - Set the field to the desired priority value (0-7). This will restrict the policy to incoming traffic that has that Inner 802.1 Priority value in the frame header. A value of 7 provides the highest priority and a value of 0 provides the lowest priority. If you do not select this option, you are effectively stating that the Inner 802.1P Priority Level is not a criterion for the Policy.

Note: Please refer to the Switch Release Notes for information on the specific QoS functions available on various platforms and combinations of hardware/firmware. Also note that if an 802.1p value is specified, a DSCP value or a ToS value may **not** be specified. This restriction does not apply to the OmniSwitch 6800 series switches.

L3 IPs

An IP Condition applies the Policy to traffic originating from, or flowing to, an IP Address/Network group. Any IP Address can be masked. Note that a Condition that specifies both a Source and Destination IP Address/Network Group will be rejected by the switch as invalid. However, if you wish to create policies for both Source and Destination traffic, you can create one policy for the Source traffic and a second policy for the Destination traffic.

Select the parameter(s) you want to configure by selecting the applicable checkbox. For Source/Destination IP Address, click on **Single** to configure a single IP Address (and **Shorthand** or **Subnet Mask**, if applicable), or click on **Group** to configure a Network Group, then enter an IP Address or select a Network Group from the drop-down menu. (You can also click the Add icon  to go to the **Groups** application and create a new Network Group.)

- **Fragment** - Select this checkbox to restrict the policy to TCP packet fragments.
- **Source IP Address/Network Group** - Configuring a Source IP Address/Network Group Condition restricts the policy to traffic that flows from this IP Address or Subnet Mask/Network Group only. If you do not select this option, you are effectively stating that the Source IP Address or Subnet

Mask/Network Group traffic is not a criterion for the policy.

- **Destination IP Address/Network Group** - Configuring a Destination IP Address/Network Group Condition restricts the policy to traffic that flows to this IP Address/Network Group only. If you do not select this option, you are effectively stating that the Destination IP Address or Subnet Mask/Network Group traffic is not a criterion for the policy.
- **Multicast IP Address Range** - Configuring a Multicast IP Address/Group Condition restricts the policy to traffic that flows to this IP Multicast Address Group only. If you do not select this option, you are effectively stating that the Destination IP Multicast Address or Subnet Mask/Group traffic is not a criterion for the policy.

Note: When configuring an IP Address Condition, you can also click either the **Shorthand Mask** or **Subnet Mask** button to configure a Subnet Mask. If you are using a Shorthand Mask, select a value from the Shorthand Mask drop-down list. If you are using a full Subnet Mask, enter the mask in the IP Subnet Mask field. Note that the * wildcard character is not allowed in IP addresses.

Important Note: When creating an IP Condition for a **NAT** Action you must specify a Network Group in the Condition. NAT will only work when both the Condition and Action specify network groups. To create a "One-to-Many" Condition and action, create a Network Group with a single entry for the Condition.

L3 DSCP/TOS

A DSCP/TOS Condition applies the Policy to incoming traffic that has a specified value in either the DSCP (Differentiated Services Code Point) byte or in the TOS (Type of Service) byte. Both DSCP and TOS are mechanisms used to convey QoS information in the IP header of frames. DSCP and TOS are mutually exclusive - you can use either DSCP or TOS but not both. Click on the applicable button (DSCP or TOS) and enter a value.

- **DSCP** - Defines the QoS treatment a frame is to receive from each network device. This is referred to as per-hop behavior. If you are using DSCP, you can define any value in the range 0 - 63 as the DSCP value in the IP header of the frame. Traffic that contains this value will match this condition.
- **TOS** - A TOS value creates a condition that applies the policy to traffic that has the specified TOS value in the IP header of frames. Enter any value from 0 - 7 to specify the value of the precedence field in the TOS byte that will match this condition. A value of 7 has the highest precedence and a value of 0 has the lowest.

Note: Please refer to the Switch Release Notes for information on the specific QoS functions available on various current platforms and combinations of hardware/firmware.

L3 TCP Flags

A TCP Flags Condition applies the Policy to traffic based on TCP values. Typically, the TCP Flags Policy Condition is used in combination with Source IP, Destination IP, Source Port, Destination Port, Source TCP Port, or Destination TCP Port conditions. Note that even though a TCP Flag condition can be used with most action parameters, it is mainly intended for ACL use. Select the parameter(s) you want to configure by selecting the applicable checkbox, then configure the parameter(s) as described below.




- **Match Established TCP Sessions**
 - **On** - Apply the policy to traffic in an established TCP session.
 - **Off** - Do not apply the policy to traffic in an established TCP session.
- **Modify The Way TCP Flags Are Matched**
 - **All** - Apply the policy to traffic that **matches all** of the TCP Flags configured in the TCP Flag

Bits fields.

- **Any** - Apply the policy to traffic that **matches any** of the TCP Flags configured in the TCP Flag Bits fields.
- **Match TCP Flags Bits**
 - **Mask Bits** - Enter one or more TCP Flags after the any or all keyword to indicate that the value of the flag bit must be set to one to qualify as a match.
 - **Match Bits** - Enter one or more TCP Flags to indicate which TCP Flags to match. If a TCP Flag is specified as part of the mask but does not have a corresponding match, a value of zero is assumed as the match value.

L4 Services

A Service Condition applies the policy to Service Protocol traffic (TCP or UDP) flowing from/to two TCP or UDP ports, or to traffic flowing from/to a TCP or UDP Service or Service Group. Select a type of Service Condition you want to configure, then configure the parameter(s) as described below.

- **Protocol Only** - Select **TCP** or **UDP** to create a condition for a Service Protocol only.
- **Port(s)** - To configure the Condition for a specific Service Port, select a **Source** and **Destination** Port from the drop-down menu to specify a specific port for the service you selected. You can also click on the Add icon  to go to the Groups application and create new Service Ports.
- **Service** - Select a Service from the drop-down menu. You can also click on the Add icon  to go to the Groups application and create a new Service.
- **Service Group** - Select a Service Group from the drop-down menu. You can also click on the Add icon  to go to the Groups application and create a new Service Group.

L7 Applications

An Application Condition is used to create a SIP Condition that applies to SIP traffic. To create a SIP Condition, select the checkbox and select a Media Type for the Condition (**Voice** / **Video** / **Other**). Selecting a Media Type, restricts the policy to that type of SIP traffic.

Application Visibility

An Application Visibility Condition applies the policy to traffic associated with a specific Application Group. Click on the **App Group** button and select an Application Group from the drop-down menu.

Note: **App Name** Conditions are **not** supported at this time.

VXLAN

A VXLAN Condition creates a VM Snooping Condition that applies to incoming VXLAN packets. VXLAN policy conditions are used to filter VXLAN packets received on VM Snooping ports. VM Snooping must be enabled on a port, and at least one parameter must be configured for a condition.

- **VXLAN VNI** - The VXLAN Network Identifier (VNI). This parameter is required to configure a VM Snooping policy condition. The VXLAN header contains the VNI that is associated with the source MAC address of the Ethernet frame that is encapsulated in a VXLAN packet. The VNI represents the VXLAN segment ID to which the packet belongs.
- **MAC Address** - The source MAC address of the VXLAN packet (source MAC address of the inner Ethernet frame of the encapsulated VXLAN packet).
- **MAC Mask** - The VXLAN Source MAC mask.
- **IP Address** - The source IP address of the packet (source IP address of the inner Ethernet frame of the encapsulated VXLAN packet). You can specify an IP v4 address/mask or an IPv6 address.

- **VXLAN Port** - The UDP destination port number. This number is found in the outer IP header of an encapsulated VXLAN packet. (Range = 0 - 65535, Default = 4789)
- **IP Protocol** - The IP protocol number (IP protocol of the inner Ethernet frame of an encapsulated VXLAN packet). (Range = 0 - 255)
- **L4 Source Port** - The Layer 4 (UDP or TCP) source port (Layer 4 port of the inner Ethernet frame of an encapsulated VXLAN packet). (Range = 0 - 65535)
- **L4 Destination Port** - The Layer 4 (UDP or TCP) destination port (Layer 4 port of the inner Ethernet frame of an encapsulated VXLAN packet). (Range = 0 - 65535)

Set Action

The Expert Mode Set Action Screen contains a list of Actions that you can configure for the Policy (e.g., QoS, NAT). A Policy Action enables you to specify the treatment traffic is to receive when it flows. This includes the priority the traffic will receive, its minimum and maximum output rates, and the values to which specified bits in the frame headers will be set upon egress from the switch. When the Conditions specified by the Policy Condition are true, traffic will flow as specified by the Policy Action.

Click on an Action to display the configuration options for the Action. (Click again on the Action to close the Action.) When you have completed all of the parameters for the Action(s), click the **Next** button at the bottom of the screen or click on Validity Period on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.

Actions

A brief description of each Action is provided below. Click the hyperlink for each Action for detailed configuration instructions.

- **QoS** - Create an Action to specify QoS actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.
- **NAT** - Create an Action to specify Network Address Translation actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.
- **PBR** - Create an Action to specify the default IP address to be used for Policy Based Routing on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.
- **TCM** - Create an Action to specify Tri-Color Marking (TCM) actions to impose on traffic that meets the configured policy condition(s). TCM provides a mechanism for policing network traffic by limiting the rate at which traffic is sent or received on a switch interface. TCM meters traffic based on user-configured packet rates and burst sizes and "marks" the metered packets as green, yellow, or red based on whether the traffic meets the configured rates. This "color marking" determines the packet's precedence when congestion occurs.
- **Ports** - Create an Action to specify QoS actions to impose on ports carrying traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.
- **SIP** - Create an Action to specify QoS actions to impose on ports carrying traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.
- **BYOD** - Create an Action to specify the BYOD Redirect Module for traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.



QoS

The QoS Policy Action option enables you to specify QoS actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

- **Disposition** - Set the Action to **Accept** or **Drop** traffic that meets the configured condition(s).
- **Quality of Service (QoS) Parameters** - Specify the QoS priority the traffic will receive if it meets the configured condition(s).
 - **Platinum** priority provides the highest quality of service (and maps to a firmware priority of 7).
 - **Gold** provides the next-highest quality of service (and maps to a firmware priority of 5).
 - **Silver** provides the next-highest quality of service (and maps to a firmware priority of 3).
 - **Bronze** provides the same quality of service as best effort (and maps to a firmware priority of 1). A separate egress queue is maintained in the hardware for traffic of each different priority.
- **Output Flow Settings**
 - **Max Output Rate (kbits/sec)** - Specify the maximum amount of traffic, in kilobits-per-second, which is guaranteed to be transmitted from the port. Even if no other traffic exists, the output will be limited to the rate specified here.
 - **Set Color of Packet** - Specify
- **Output Mapping**
 - **802.1p Priority Level** - If you want outgoing packets tagged with an 802.1p priority level, set the **802.1p Priority Level** field to any value between 0 to 7 to specify the desired outgoing 802.1p priority for the traffic. A value of **7** indicates the highest priority and a value of 0 indicates the lowest priority. Note that for ports that are configured for 802.1q, this value is used in the 802.1q header and indicates the outgoing priority of the frame. When a frame is de-queued for transmission, it is assigned the priority of the queue and mapped to the outgoing 802.1p priority. This priority is combined with the VLAN group ID to create the 802.1p/q header for transmission. Note that if traffic matches the criteria specified by the policy condition, but the outgoing port does not support 802.1p tagging, the policy action will fail.
 - **DSCP/TOS** - Enable/Disable DSCP/TOS Precedence. The TOS byte is defined in RFC 791. This byte contains two fields. The precedence field is the three high-order bits (0-2) and is used to indicate the priority for the frame. The type of service field (bits 3-6) defines the throughput, delay, reliability, or cost for the frame; however, in practice these bits are not used. If you enable the **TOS Precedence** radio button, set the associated field to any value from 0-7 to specify the value that will be inserted into the precedence field of the TOS byte upon egress from the switch. A value of 7 has the highest precedence and a value of 0 has the lowest precedence. Note that you can enable **either** the DSCP or the TOS Precedence radio button to specify the mechanism you want to use (if any) to convey QoS information in the IP header of frames. DSCP and TOS are mutually exclusive. You can use either DSCP or TOS, but not both.

NAT

The NAT Policy Action option enables you to specify Network Address Translation actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

- **Source Rewrite IP Address** - To include Source Rewrite IP in the NAT Policy condition, select Network Group to be used for policy condition from the **Source Rewrite IP Address** drop-down menu. You can also click on the Add icon  to go to the Network Groups Screen and create a
- Network Group. **Destination Rewrite IP Address** - To include Destination Rewrite IP in the NAT Policy condition, select Network Group to be used for policy condition from the **Destination Rewrite IP Address** drop-down menu. You can also click on the Add icon  to go to the Network Groups Screen and create a Network Group.

Note: Remember, when creating a condition (e.g., MAC, IP) for a NAT action you must specify a group in the condition. NAT will only work when both the condition and the action specify groups. To create a "one-to-many" condition and action, create a group with a single entry for the condition. Also note that the NAT Policy Action is **not supported on OS6860, OS6900, or OS10K** Switches.

PBR

The PBR Policy Action option enables you to specify the default IP address to be used for Policy Based Routing on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

- **Permanent Gateway IP** - To set a Permanent Gateway IP address for traffic that meets the condition(s), enter the default IP address in the **PBR Permanent Gateway IP Address** field.
- **Alternate Gateway IP** - To specify an alternate IP address for traffic that meets the policy condition(s), enter the alternate IP address in the **PBR Alternate Gateway IP Address** field.

Note : The OmniSwitch 6800/7000/8000/9000 series switches support the 802.1 priority, DSCP, and TOS. However, 6600 series switches do not. Please refer to the switch Release Notes for information on the specific QoS functions available on various current platforms and combinations of hardware/firmware.

TCM

The TCM Policy Action option enables you to specify Tri-Color Marking (TCM) actions action to impose on traffic that meets the configured policy condition(s). TCM provides a mechanism for policing network traffic by limiting the rate at which traffic is sent or received on a switch interface. TCM meters traffic based on user- configured packet rates and burst sizes and "marks" the metered packets as green, yellow, or red based on whether the traffic meets the configured rates. This "color marking" determines the packet's precedence when congestion occurs.

- **Committed Traffic Policing**
 - **Committed Information Rate** - The maximum amount of bandwidth, in bits-per-second, for all traffic that ingresses on the port.
 - **Committed Burst Size** - The maximum burst size, in bits-per-second, for all traffic that ingresses on the port.
- **Peak Traffic Policing**
 - **Peak Information Rate** - The maximum amount of bandwidth, in bits-per-second, for all traffic that ingresses on the port.
 - **Peak Burst Size** - The maximum burst size, in bits-per-second, for all traffic that ingresses on the port.

Ports

The Ports Policy Action option enables you to specify QoS actions to impose on ports carrying traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action. Select the applicable checkbox as described below and configure the mirroring slot/port.

Slot/Port Mirroring

The Slot/Port Mirroring fields are used to mirror ingress, egress, or both ingress and egress packets that match the policy condition to the specified port. Note that only one MTP session is supported at any given time. As a result, all mirroring policies should specify the same MTP port.

- **Slot/Port Mirroring** - For a non-Virtual Chassis (VC) Switch, enter the mirroring **Slot** and **Port** number and select the **Traffic Direction** from the drop-down menu (Ingress, Egress, Ingress/Egress).
- **Chassis/Slot/Port Mirroring for VC Devices** - For a VC Switch, enter mirroring **Chassis ID**, **Slot**, and **Port**, and select the **Traffic Direction** from the drop-down menu (Ingress, Egress, Ingress/Egress).

Slot/Port Redirection

The Slot/Port Redirection fields are used to redirect all traffic (flooded, bridged, routed, and multicast) matching a the policy condition to the specified port instead of the port to which the traffic was originally destined. Note that when redirecting routed traffic from VLAN A to VLAN B, the redirect port must belong to VLAN B (tagged or default VLAN). Also, routed packets (from VLAN A to VLAN B) are not modified after they are redirected; the source and MAC address remain the same. In addition, if the redirect port is tagged, the redirected packets will have a tag from the ingress VLAN A.

- **Slot/Port Redirection** - For a non-Virtual Chassis (VC) Switch, enter the **Slot** and **Port** number to which you want the traffic re-directed.
- **Chassis/Slot/Port Redirection for VC Devices** - For a VC Switch, enter the **Chassis ID** and **Slot/Port** or **Link Aggregate**, for the slot/port or link aggregate to which you want the traffic re-directed.

Port Disable Rule Match

Enable this option to administratively disable the source port of the traffic matching the policy condition(s).

SIP

The SIP Policy Action option enables you to specify QoS actions to impose on ports carrying traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

- **RTCP Monitoring** - Enables/Disables monitoring of RTCP Marked traffic. If Enabled, traffic meeting the configured condition(s) will be subjected to RTCP Monitoring.
- **RTCP DSCP** - The **RTCP DSCP** number is used as a prioritizing rate number for SIP PDUs. To apply an RTCP-DSCP number to traffic meeting the configured condition(s), enter a value (Range = 0 to 63, Default = 46).
- **Trust DSCP** - If Enabled, traffic meeting the configured condition(s) will have the "Trust DHCP" function applied.

Note: The SIP feature is only supported on the following devices running AOS 6.4.5.R02 and later: 6850E (C24/24x/48/48X, P24/24X/48/48X,U24X), 6855 (U24x), 9700E (C-24/48, P24, U2/6/12/24), 9800E (C24/48, P24, U2/6/12/24).

BYOD

The BYOD Policy Action option enables you to specify the BYOD Redirect Module for traffic that meets the configured policy condition(s) (None, QMR, Captive Portal, Unauthorized BYOD).

Validity Period

The Expert Mode Validity Period Screen enables you to add a validity period to a condition by specifying the time periods when the policy is active and enforced. Four pre-configured policy validity periods are provided in the drop-down list in the Policy Validity Periods pane. They are AllTheTime, Weekdays, Weekends, and WorkingDay . You can also create Custom validity periods that are enforced during a specific timeframe.

When you have completed all of the parameters, click the Next button at the bottom of the screen or click on Review on the left side of the screen to move to the next step. If necessary, you can also click the Back button to return to the screen.

Note: The pre-configured validity period **AllTheTime** is the default and is automatically assigned to the condition when the **Ignore Validity Period in defining Policy Condition** checkbox is checked.

Review

The Expert Mode Review Screen is used to review the Policy configuration before saving the Policy to the LDAP Server. After reviewing the Policy, click the **Create** button to save the policy to the LDAP Server. You can also click the **Back** button to return to a previous screen.

Applications


The PolicyView Application Visibility option enables you to create Application Visibility Policies and Policy Lists for Application traffic flows. When you select "Create Policy for Applications" from the main PolicyView screen, the Application Visibility Policies Screen appears. You can configure Application Visibility Policies from this screen or click the Application Visibility Policy List link on the left side of the screen to configure Application Visibility Policy Lists. Click on the Application Visibility Policies link to return to the Application Visibility Policies Screen.

Application Visibility Policies

The PolicyView Application Visibility Policies Screen displays and is used to create, edit, delete, and view Application Visibility Policies. The Policies are created using a wizard that guides you through each of the steps needed to create the Policy and apply the Policy to switches in the network. As shown below, all currently-configured Policies are listed in the Existing Policies Table. You can view basic information about a Policy in the table or click on a specific policy to view more detailed information.

Note: Application Visibility Policies are configured using Application Groups configured as part a Signature Profile created in the Application Visibility Application. The drop-down menu on the Set Condition Screen in the wizard is populated with the Application Groups contained in the Signature Profile for the selected switch(es). You must configure and apply a Signature Profile to a switch **before** creating an Application Visibility Policy.

Creating an Application Visibility Policy

Application Visibility Policies are created using a wizard that guides you through each of the steps needed to create the policy and apply the policy to switches in the network. To create an Application Visibility Policy, click on the Create icon . The wizard will then guide you through the following screens:

- Configuration - Basic policy configuration (e.g., Policy Name, Precedence)
- Device Selection - Specify the devices to which you will apply the policy
- Set Condition - Specify the conditions that must be true before traffic will be allowed to flow.
- Set Action - Specify parameters for the traffic that will flow.
- Validity Period - Specify the time period for the policy to be in effect.
- Review - Review the policy details before applying the policy to the network.


Applying an Application Visibility Policy to the Network

After configuring and saving a policy(ies), you must go to the PolicyView Expert Screen to notify the switches. The policy(ies) you created will appear in the Existing Policies Table. Click on the Notify All button to notify all switches in the network; or click on the Select Devices button to notify specific switches. When you click on the Notify All button, all of the policies listed in the Existing Policies Table are applied to all of the switches configured for each policy. To apply the policy(ies) only to certain switches within the configured group of switches, click on the Select Devices button and use the Device Selection Wizard to select specific devices (select the device(s) and click OK; then click the Notify Selected button).

After notifying the switches, you can view the status of the re-cache operation, by clicking on the Status button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the Audit application, including an indication of any error that may have occurred. Note that the re-cache operation for each switch occurs in a separate thread and may take some time. Any errors that occur will also be reported in the server.txt file, in the Audit application.

Note: When you notify network switches, all QoS-enabled switches flush their policy tables and reload policies from the LDAP repository, which is very expensive in terms of switch resources and time. It is recommended that you review all policies that you have created and apply them at the same time to minimize switch downtime.

Editing an Application Visibility Policy

To edit a policy, select the policy in the Existing Application Visibility Policies Table and click on the Edit icon . Use the wizard to make any edits. When you are done, apply the edited policy to the network.

Deleting an Application Visibility Policy

To delete a policy(ies), select the policy(ies) in the Existing Application Visibility Policies Table and click on the Delete icon, then click **OK** at the confirmation prompt.

Policy Information

The Existing Application Visibility Policies Table displays basic information for all configured Policies. You can also click on a policy to view detailed information about the Policy (e.g., Condition, Action).

Basic Information

- **Policy Name** - The name of the Policy.
- **Scope** - The scope of the Policy (e.g., IP Filtering, Provisioned).
- **Precedence** - The Precedence value of the Policy (0 - 65535).
- **Status** - Indicates whether or not the Policy has been saved to the LDAP Server.
- **Enable** - Indicates whether or not the Policy is enabled.
- **Save** - Indicates whether or not the rule will be recorded during a snapshot command.
- **Log Matches** - Indicates whether or not matches to this rule are logged in the QoS Log.
- **Reflexive** - Indicates whether or not the Policy is reflexive. Reflexive Policy Rules allow specific return connections that would normally be denied (Yes/No).
- **Default List** - Indicates whether or not the Policy is saved to the Default Policy List. By default, a Policy Rule is added to this list when it is created. A Policy Rule remains a member of the Default List even when it is subsequently assigned to additional Policy Lists.
- **SLA Policy Trap** - Indicates whether or not an SLA Policy Trap is configured for the policy.

Detailed Information

- **Policy Rule** - The name of the Policy Rule and the Policy Rules configured for the Policy.
- **Policy Condition** - The Policy condition information (e.g., IP Policy condition would display the Source/Destination/Multicast IP address of the condition).
- **Policy Action** - The Policy action to take if the traffic matches the Policy condition (e.g., QoS Accept/Drop)
- **Policy Validity Period** - The configured validity period for the Policy.
- **Policy Roles** - The switches to which the Policy has been assigned.

Config for Policy

The Application Visibility Policy Config for Policy Screen is used to configure basic Policy parameters. When you have completed all of the parameters, click the **Next** button at the bottom of the screen or click on Device Selection on the left side of the screen to move to the next step.

- **Name** - The Policy name.
- **Precedence** - The Policy precedence. By default, the precedence field is pre-filled with the lowest unused precedence value (Range = 0 - 65535).

Click on **Show Advanced Options** to display and configure the options below. By default, these options are set to **Ignore**.

- **Default List** - Adds the rule to the QoS Default Policy List.
- **Enable** - Enables the policy.
- **Save** - Marks the policy rule so that it may be captured as part of the switch configuration.
- **Log Matches** - Configures the switch to log messages about specific flows coming into the switch that match this policy rule.
- **Send Trap** - Enables traps for the Policy.
- **Reflexive** - Enables support for the Reflexive for the policy. Reflexive policies allow specific return connections that would normally be denied.

Device Selection

The Application Visibility Policy Device Selection Screen is used to select the switches to which you want to apply the Policy. Select an option (Use Switch Picker/Use Topology), and select the device(s). Only switches with an assigned Signature Profile are available. Click on the **Next** button at the bottom of the screen or click on Set Condition on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen and add/delete devices.

Set Condition

The Application Visibility Policy Set Condition Screen is used to select the Application Group/Application for the Policy. When you create a Condition, traffic matching this Application Group/Application will be subject to the Action that you configure. Click on the **App Group** button and select a Group for the condition.

Note that the drop-down menus are populated with the Application Groups/Applications contained in the Signature Profile for the selected switch. If you select multiple switches, only those Application Groups/Applications common to all switches will be displayed. Also note that the **App Name** button will not be displayed if you select any OS6900 Switches, as this option is not offered for these devices. If all of the selected switches are OS6860 devices, both the **App Group** and **App Name** buttons are displayed.

When you have selected the Application Group/Application, click the **Next** button at the bottom of the screen or click on Set Action on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.

Set Action

The Application Visibility Policy Set Action Screen contains a list of Actions that you can configure for the Policy (e.g., QoS, NAT). A Policy Action enables you to specify the treatment traffic is to receive when it flows. This includes the priority the traffic will receive, its minimum and maximum output rates, and the values to which specified bits in the frame headers will be set upon egress from the switch. When the Conditions specified by the Policy Condition are true, traffic will flow as specified by the Policy Action.

Click on an Action to display the configuration options for the Action. (Click again on the Action to close the Action.) When you have completed all of the parameters for the Action(s), click the **Next** button at the bottom of the screen or click on Validity Period on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.

Actions

A brief description of each Action is provided below. Click the hyperlink for each Action for detailed configuration instructions.

- **QoS** - Create an Action to specify QoS actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.
- **NAT** - Create an Action to specify Network Address Translation actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.
- **PBR** - Create an Action to specify the default IP address to be used for Policy Based Routing on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.
- **TCM** - Create an Action to specify Tri-Color Marking (TCM) actions to impose on traffic that meets the configured policy condition(s). TCM provides a mechanism for policing network traffic by limiting the rate at which traffic is sent or received on a switch interface. TCM meters traffic based on user-

configured packet rates and burst sizes and "marks" the metered packets as green, yellow, or red based on whether the traffic meets the configured rates. This "color marking" determines the packet's precedence when congestion occurs.

- **Ports** - Create an Action to specify QoS actions to impose on ports carrying traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.
- **SIP** - Create an Action to specify QoS actions to impose on ports carrying traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

QoS



The QoS Policy Action option enables you to specify QoS actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

- **Disposition** - Set the Action to **Accept** or **Drop** traffic that meets the configured condition(s).
- **Quality of Service (QoS) Parameters** - Specify the QoS priority the traffic will receive if it meets the configured condition(s).
 - **Platinum** priority provides the highest quality of service (and maps to a firmware priority of 7).
 - **Gold** provides the next-highest quality of service (and maps to a firmware priority of 5).
 - **Silver** provides the next-highest quality of service (and maps to a firmware priority of 3).
 - **Bronze** provides the same quality of service as best effort (and maps to a firmware priority of 1). A separate egress queue is maintained in the hardware for traffic of each different priority.
- **Output Flow Settings**
 - **Max Output Rate (kbits/sec)** - Specify the maximum amount of traffic, in kilobits-per-second, which is guaranteed to be transmitted from the port. Even if no other traffic exists, the output will be limited to the rate specified here.
 - **Set Color of Packet** - Specify
- **Output Mapping**
 - **802.1p Priority Level** - If you want outgoing packets tagged with an 802.1p priority level, set the **802.1p Priority Level** field to any value between 0 to 7 to specify the desired outgoing 802.1p priority for the traffic. A value of **7** indicates the highest priority and a value of 0 indicates the lowest priority. Note that for ports that are configured for 802.1q, this value is used in the 802.1q header and indicates the outgoing priority of the frame. When a frame is de-queued for transmission, it is assigned the priority of the queue and mapped to the outgoing 802.1p priority. This priority is combined with the VLAN group ID to create the 802.1p/q header for transmission. Note that if traffic matches the criteria specified by the policy condition, but the outgoing port does not support 802.1p tagging, the policy action will fail.
 - **DSCP/TOS** - Enable/Disable DSCP/TOS Precedence. The TOS byte is defined in RFC 791. This byte contains two fields. The precedence field is the three high-order bits (0-2) and is used to indicate the priority for the frame. The type of service field (bits 3-6) defines the throughput, delay, reliability, or cost for the frame; however, in practice these bits are not used. If you enable the **TOS Precedence** radio button, set the associated field to any value from 0-7 to specify the value that will be inserted into the precedence field of the TOS byte upon egress from the switch. A value of 7 has the highest precedence and a value of 0 has the lowest precedence. Note that you can enable **either** the DSCP or the TOS Precedence radio button to specify the mechanism you want to use (if any) to convey QoS information in the IP header of

frames. DSCP and TOS are mutually exclusive. You can use either DSCP or TOS, but not both.

NAT

The NAT Policy Action option enables you to specify Network Address Translation actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

- **Source Rewrite IP Address** - To include Source Rewrite IP in the NAT Policy condition, select Network Group to be used for policy condition from the **Source Rewrite IP Address** drop-down menu. You can also click on the Add icon  to go to the Network Groups Screen and create a
- Network Group. **Destination Rewrite IP Address** - To include Destination Rewrite IP in the NAT Policy condition, select Network Group to be used for policy condition from the **Destination Rewrite IP Address** drop-down menu. You can also click on the Add icon  to go to the Network Groups Screen and create a Network Group.

Note: Remember, when creating a condition (e.g., MAC, IP) for a NAT action you must specify a group in the condition. NAT will only work when both the condition and the action specify groups. To create a "one-to-many" condition and action, create a group with a single entry for the condition. Also note that the NAT Policy Action is **not supported on OS6860, OS6900, or OS10K** Switches.

PBR

The PBR Policy Action option enables you to specify the default IP address to be used for Policy Based Routing on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

- **Permanent Gateway IP** - To set a Permanent Gateway IP address for traffic that meets the condition(s), enter the default IP address in the **PBR Permanent Gateway IP Address** field.
- **Alternate Gateway IP** - To specify an alternate IP address for traffic that meets the policy condition(s), enter the alternate IP address in the **PBR Alternate Gateway IP Address** field. (**Not supported on OS6860 or OS6900** Switches.)

Note: The OmniSwitch 6800/7000/8000/9000 series switches support the 802.1 priority, DSCP, and TOS. However, 6600 series switches do not. Please refer to the switch Release Notes for information on the specific QoS functions available on various current platforms and combinations of hardware/firmware.

TCM

The TCM Policy Action option enables you to specify Tri-Color Marking (TCM) actions action to impose on traffic that meets the configured policy condition(s). TCM provides a mechanism for policing network traffic by limiting the rate at which traffic is sent or received on a switch interface. TCM meters traffic based on user-configured packet rates and burst sizes and "marks" the metered packets as green, yellow, or red based on whether the traffic meets the configured rates. This "color marking" determines the packet's precedence when congestion occurs.

- **Committed Traffic Policing**
 - **Committed Information Rate** - The maximum amount of bandwidth, in bits-per-second, for all traffic that ingresses on the port.
 - **Committed Burst Size** - The maximum burst size, in bits-per-second, for all traffic that

ingresses on the port.

- **Peak Traffic Policing**

- **Peak Information Rate** - The maximum amount of bandwidth, in bits-per-second, for all traffic that ingresses on the port.
- **Peak Burst Size** - The maximum burst size, in bits-per-second, for all traffic that ingresses on the port.

Ports

The Ports Policy Action option enables you to specify QoS actions to impose on ports carrying traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action. Select the applicable checkbox as described below and configure the mirroring slot/port.

Slot/Port Mirroring

The Slot/Port Mirroring fields are used to mirror ingress, egress, or both ingress and egress packets that match the policy condition to the specified port. Note that only one MTP session is supported at any given time. As a result, all mirroring policies should specify the same MTP port.

- **Slot/Port Mirroring** - For a non-Virtual Chassis (VC) Switch, enter the mirroring **Slot** and **Port** number and select the **Traffic Direction** from the drop-down menu (Ingress, Egress, Ingress/Egress).
- **Chassis/Slot/Port Mirroring for VC Devices** - For a VC Switch, enter mirroring **Chassis ID**, **Slot**, and **Port**, and select the **Traffic Direction** from the drop-down menu (Ingress, Egress, Ingress/Egress).

Slot/Port Redirection

The Slot/Port Redirection fields are used to redirect all traffic (flooded, bridged, routed, and multicast) matching a the policy condition to the specified port instead of the port to which the traffic was originally destined. Note that when redirecting routed traffic from VLAN A to VLAN B, the redirect port must belong to VLAN B (tagged or default VLAN). Also, routed packets (from VLAN A to VLAN B) are not modified after they are redirected; the source and MAC address remain the same. In addition, if the redirect port is tagged, the redirected packets will have a tag from the ingress VLAN A.

- **Slot/Port Redirection** - For a non-Virtual Chassis (VC) Switch, enter the **Slot** and **Port** number to which you want the traffic re-directed.
- **Chassis/Slot/Port Redirection for VC Devices** - For a VC Switch, enter the **Chassis ID** and **Slot/Port** or **Link Aggregate**, for the slot/port or link aggregate to which you want the traffic re-directed.

Port Disable Rule Match

Enable this option to administratively disable the source port of the traffic matching the policy condition(s).

SIP

The SIP Policy Action option enables you to specify QoS actions to impose on ports carrying traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

- **RTCP Monitoring** - Enables/Disables monitoring of RTCP Marked traffic. If enabled, traffic meeting the configured condition(s) will be subjected to RTCP Monitoring.
- **RTCP DSCP** - The **RTCP DSCP** number is used as a prioritizing rate number for SIP PDUs. To apply an RTCP-DSCP number to traffic meeting the configured condition(s), enter a value (Range = 0 to 63, Default = 46).
- **Trust DSCP** - If enabled, traffic meeting the configured condition(s) will have the "Trust DSCP" function applied.

Note: The SIP feature is only supported on the following devices running AOS 6.4.5.R02 and later: 6850E (C24/24x/48/48X, P24/24X/48/48X,U24X), 6855 (U24x), 9700E (C-24/48, P24, U2/6/12/24), 9800E (C24/48, P24, U2/6/12/24).

Validity Period

The Application Visibility Policy Validity Period Screen enables you to add a validity period to a condition by specifying the time periods when the policy is active and enforced. Four pre-configured policy validity periods are provided in the drop-down list in the **Policy Validity Periods** pane. They are **AllTheTime**, **Weekdays**, **Weekends**, and **WorkingDay**. You can also create **Custom** validity periods that are enforced during a specific timeframe.

When you have completed all of the parameters, click the **Next** button at the bottom of the screen or click on Review on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.

Note: The pre-configured validity period **AllTheTime** is the default and is automatically assigned to the condition when the **Ignore Validity Period in defining Policy Condition** checkbox is checked.

Review

The Application Visibility Policy Review Screen is used to review the Policy configuration before saving the Policy to the LDAP Server. After reviewing the Policy, click the **Create** button to save the policy to the LDAP Server. You can also click the **Back** button to return to a previous screen.

Application Visibility Policy Lists

The PolicyView Application Visibility Policy Lists Screen displays all configured Application Visibility Policy Lists, including the Policy Rules included in each list, and is used to create, edit, and delete Policy Lists. You can click on a Policy List to display the Policy Rules contained in the list. A Policy List is a set of Policies that are grouped together and can be assigned to switches as a group. The QoS/ACL Policies that you add to a Policy List can be defined using the PolicyView Application. You can also include Resources or Resource Groups in a Policy using the "Resources - Add to Policy" Screen.

Note: Application Visibility is supported on OS10K Switches (AOS 7.3.4.R01 and later), OS6900 Switches (AOS 7.3.4.R02 and later) and OS6860/6860E Switches (AOS 8.2.1.R01 and later). It is also supported in a virtual chassis of OS6860/OS6860E Switches where at least one OS6860E is present.



Note: Application Visibility Policy Lists can only be configured on OS6900 Switches (AOS 7.3.3.R01 and later). The Application Enforcement feature in the Application Visibility application is only supported on OS6860 Switches (AOS 8.1.1.R01 and later). However, Application Enforcement can be configured on OS6900 Switches by creating Application Policy Lists, and using them to create a QoS Application Policies for application enforcement.

Policy List Information



The following information is displayed for each Policy contained in the Policy List. (Click on a Policy List to display the Policies contained in the list.)

- **Name** - The name of the Policy.
- **Condition** - The Policy Condition information (e.g., IP Policy condition would display the Source/Destination/Multicast IP address of the condition).
- **Action** - The Policy Action to take if the traffic matches the Policy condition (e.g., QoS Accept/Drop)
- **Precedence** - The Precedence value of the Policy (0 - 65535).
- **Validity Period** - The configured validity period for the Policy.

Creating a Policy List


Click on the Create icon . The Create Policy List Screen appears. Enter a **Name** for the Policy List and select the Policies you want to include in the list from the **Add Policies** drop-down menu. (All of the currently- configured Policies appear in the list. You can also click the Add icon  to go to Application Visibility Policies Screen and create a new policy(ies) to add to the list.) When you select a Policy from the drop-down menu, the Policy will appear in a table below, so you can review the Policy and modify the Precedence value, if needed. When you are finished reviewing the Policy(ies), click the **Create** button. The new Policy List will appear on the Policy Lists Screen.

Editing a Policy List

You can edit the Policies included in a Policy List or edit the Precedence value of any Policy in the list. Select a Policy List and click on the Edit icon . The Edit Policy List Screen appears. Click on the **Add Policies** drop-down menu. (All of the currently-configured Policies appear in the list. You can also click the Add icon  to go to Application Visibility Policies Screen and create a new policy(ies) to add to the list.) Select/unselect Policies to add/remove them from the Policy List. When you are finished editing the Policy, click the **Update** button. The updated Policy List will appear on the Policy Lists Screen.

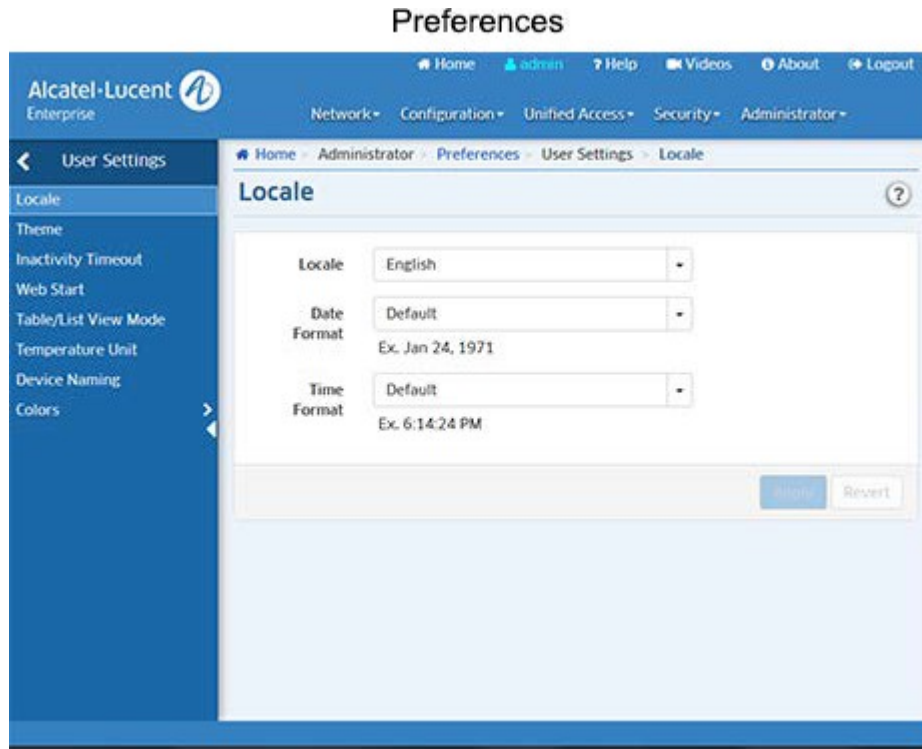
Note: Adding/deleting a policy to/from a policy list will automatically update any switch roles that contain this policy list with the updated policies.

Deleting a Policy List

To delete a Policy List(s), select the list(s), click on the Delete icon , then click **OK** at the confirmation prompt. Note that you cannot delete a Policy List that is associated with a User Network Profile (UNP). To delete the list, you must first remove it from the UNP.

18.0 Preferences

The Preferences Application is used to set OmniVista preferences for the web GUI. All Preferences have appropriate default values, so there is no need to change Preference settings unless you wish to. When a Preference is changed, the change takes effect immediately.



The following preferences can be configured. Any user can update their User settings; however a user must be assigned to the Account Admin Role to configure System settings.

User Settings - These settings can be configured for each user.

- **Locale** - Used to set a system-wide language, and time/date format.
- **Theme** - Used to set the color scheme and look of OmniVista.
- **Inactivity Timeout** - Used to set the Inactivity Timer. If there is no user activity within this timeframe, the user is logged off.
- **Table/List View Mode** - Used to set the default display layout for all table/list screens in OmniVista.
- **Temperature Unit** - Used to set the temperature unit that will be displayed, when applicable, in OmniVista (e.g., Centigrade or Fahrenheit).
- **Device Naming** - Used to specify how devices are identified and displayed in OmniVista (e.g., IP address, Device Name, DNS Name).
- **Colors** - Used to configure the colors displayed in Dashboard Widgets for Network Status, Alarms, Quarantine Manager, and ProActive Lifecycle Management.

System Settings - These settings are system-wide settings that are configured for all users.

- **Branding** - Used to change the logo displayed on the OmniVista user interface and the logo displayed on reports created in the Report application.
- **Proxy** - Used to configure a Proxy for the OmniVista Client.
- **ProActive Lifecycle Management** - Used to enable/disable ProActive Lifecycle Management and manually upload information.

- **Videos** - Used to specify the Alcatel-Lucent Enterprise YouTube Demo Playlist that will play when the "Videos" link at the top of the OmniVista screen is clicked.
- **Email** - Used to specify the Simple Mail Transfer Protocol (SMTP) mail server that you want to use to send e-mails generated by OmniVista.
- **Install Zulu CEK** - Used to install the Zulu Cryptography Extension Kit (CEK).

Locale

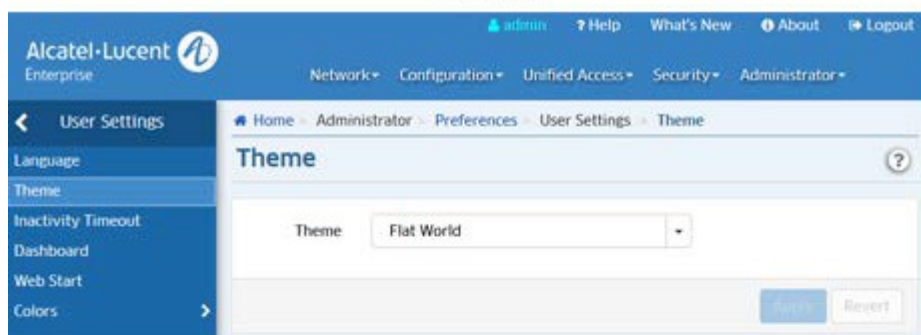
The Preferences Locale Screen is used to set a system-wide language, and time/date format. When you have configured a parameter, click the **Apply** button. The changes take effect immediately.

- **Language** - The language that will display (Default = English)
- **Date Format** - The date format that will display:
 - **Full** - Day and Date (e.g., Wednesday, August 20, 2014)
 - **Medium** - Abbreviated Month, Day, Year (e.g., Aug 20, 2014)
 - **Long** - Full Month, Day Year (e.g., August 20, 2014)
 - **Default** - Medium format (abbreviated Month, Day, Year (e.g., Aug 20, 2014)
- **Time Format** - The time format that will display:
 - **Short** - Hours:Minutes - AM-PM (e.g., 3:15 PM)
 - **Full** - Hours:Minutes:Seconds - AM-PM - Timezone (e.g., 3:15:50 PM PST)
 - **Medium** - Hours:Minutes:Seconds - AM-PM (e.g, 3:15:50 PM)
 - **Long** - Hours:Minutes:Seconds - AM-PM - Timezone (e.g., 3:15:50 PM PST)
 - **Default** - Medium format (Hours:Minutes:Seconds - AM-PM (e.g., 3:15:50 PM)

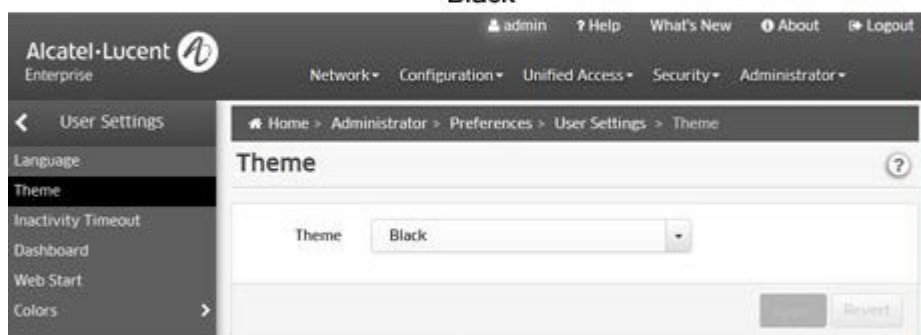
Theme

The Preferences Theme Screen is used to set the color scheme and look of OmniVista. When you have selected a theme, click the **Apply** button. The change takes effect immediately. The available themes are shown below.

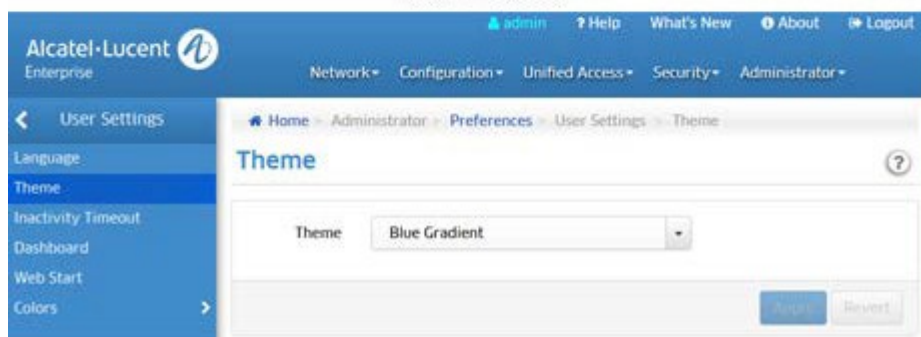
Flat World



Black



Blue Gradient



Inactivity Timeout

The Preferences Inactivity Timeout Screen is used to set the Inactivity Timer. If there is no user activity within this timeframe, the user is logged off. Enter a time, in minutes. You can also use the **+/-** symbols to increase/decrease the time by 5 minute increments, or use the slider for larger increments. When you have configured the time, click the **Apply** button. The change takes effect immediately. (Range = 15 - 525,600 (1 Year), Default = 15).

Table/List View Mode

The Preferences Table/List View Mode Screen is used to set the default display layout for all table/list screens in OmniVista. Note that you can always change the view on a table/list screen using the view option icons at the top of the screen. Select a display mode (Table or List) and click on the **Apply** button. The changes take effect immediately.

Temperature Unit

The Preferences Temperature Unit Screen is used to set the temperature unit that will be displayed, when applicable, in OmniVista. From the **Temperature Unit** drop-down menu, select "Centigrade" or "Fahrenheit" and click on the **Apply** button. The change takes effect immediately.

Device Naming Pattern

The Preferences Device Naming Pattern Screen is used to specify how devices are identified and displayed in OmniVista (e.g., IP address, Device Name, DNS Name). This preference sets the device naming style for all applications within OmniVista. Select an option from the **Device Naming Pattern** drop-down menu and click on the **Apply** button. The change takes effect immediately.

Network Status Color Preferences

The Preferences Network Status Color Preferences Screen is used to set the colors that display in the Topology application and Network Status widget on the Dashboard. You can change colors by clicking on a color and using the color picker; or you can enter a color using RGB Hexadecimal format (e.g., #2ca02c). When you have configured a setting, click the **Apply** button. The change takes effect immediately.

Alarms Color Preferences

The Preferences Alarms Color Preferences Screen is used to set the colors for alarm notifications that display in the Notifications application (and in the Topology application). You can change colors by clicking on a color and using the color picker; or you can enter a color using RGB Hexadecimal format (e.g., #2ca02c). When you have configured a setting, click the **Apply** button. The change takes effect immediately.

Quarantine Manager Color Preferences

The Preferences Quarantine Manager Color Preferences Screen is used to set the notification colors that display in the Quarantine Manager widget on the Dashboard. You can change colors by clicking on a color and using the color picker; or you can enter a color using RGB Hexadecimal format (e.g., #2ca02c). When you have configured a setting, click the **Apply** button. The change takes effect immediately.

ProActive Lifecycle Management Color Preferences

The Preferences ProActive Lifecycle Management Color Preferences Screen is used to set the colors that display for each status in the ProActive Lifecycle Management widget on the OmniVista Dashboard. Each status level (e.g., Supported, Not Supported) in the widget pie charts is also assigned a level number. For example, Level 0 corresponds to "Supported", Level 1 corresponds to "Support to End", and Level 2 corresponds to "Not Supported. (You can hover over a section of a pie chart to see the Level number that corresponds to a status level.) To change the color that will be displayed for a status level, change the corresponding Level number color on the Preferences Screen.

To configure a color, click on the color box next to a level, then use the color picker or enter a color using RGB Hexadecimal format (e.g., #2ca02c) to change the color. When you have configured a color, click the **Apply** button. The change takes effect immediately. Click on the **Restore Default Colors** button to return all colors to their default settings.

Branding

The Preferences Branding Screen is used to change the logo displayed on the OmniVista user interface and the logo displayed on reports created in the Report application. By default, the Alcatel-Lucent Enterprise logo is displayed. However, you can upload a custom logo to be displayed. To upload a custom logo, click on the "Upload Custom Logo" link to locate and upload the logo. The logo must conform to the size and dimensions shown. You can also select a dark or light background for the logo. Click on the applicable option to view how it will be displayed.

After uploading the file and selecting a background, click on the **Apply** button. The new logo will immediately appear in the upper left corner of the screen, replacing the Alcatel-Lucent Enterprise logo; and will appear on any reports you create. At any time, you can click on the "Use Default Logo" link and click the **Apply** button to return to the default Alcatel-Lucent Enterprise logo.

Proxy

The Preferences Proxy Screen is used to configure a web proxy for the Asset Management (Call Home) and Application Visibility Signature File Update Features. To configure a proxy, complete the fields and click on the **Apply** button. The change takes effect immediately.

ProActive Lifecycle Management

The Preferences ProActive Lifecycle Management (PALM) Screen is used to configure PALM preferences. The PALM Feature periodically gathers detailed inventory information for all discovered devices on your network (e.g., device name, MAC address, AOS version, NIs, power supplies) and uploads the information to the PALM web portal. The information can assist Alcatel-Lucent Enterprise in helping you manage your inventory and your network; and is also available to you through a widget that can be displayed on the OmniVista 2500 NMS Dashboard for easy reference.

After updating any preferences, click on the **Apply** button. The change(s) take effect immediately.

- **OV ID** - The ID displayed is the PALM ID automatically assigned to your OmniVista Server at installation and is used to identify your system. It is not configurable.
- **Backend URL** - This is the URL for the PALM web portal. There is no need to change the URL unless directed to do so by Alcatel-Lucent Enterprise.
- **ProActive Lifecycle Management** - Enables/Disables the PALM Feature. When you install OmniVista 2500 NMS, the PALM option is selected by default on the License Agreement Screen. If you accept the default, PALM is automatically enabled following installation. If you opt out of the feature at installation and need to enable it later, click on the **ProActive Lifecycle Management** slider to select **Enabled**.

Accept the License Agreement and click the **Accept** Button. To disable the feature, click on the **ProActive Lifecycle Management** slider to select **Disabled**. If enabled, you can click on the **Upload Now** button to perform an immediate upload of PALM data.

- **Inventory Status** - Displays the time and date of the last successful upload of PALM data, and the time and date of the next scheduled upload. Until the first successful upload, the field will display "Never".
After an initial upload at installation, updated information is sent to the web portal every two (2) weeks.
- **PALM Widget Status** - Displays the time and date of the last successful sync with the ProActive Lifecycle Management widget displayed on the Dashboard, and the time and date of the next scheduled sync. Until the first successful sync, the field will display "Never".

Note: You will be prompted to enable the PALM Feature whenever you add/relicense an OmniVista Core License in the License application. If necessary, you can always enable the feature on this screen.

Videos

The Preferences Videos Screen is used to specify the Alcatel-Lucent Enterprise YouTube Demo Playlist that will play when the "Videos" link at the top of the screen is clicked. Enter the **YouTube Playlist ID**. (You can view the widow in the **Preview** window for confirmation.) Click on the **Apply** button to set the playlist.

Email

The Preferences Email Screen is used to specify the Simple Mail Transfer Protocol (SMTP) mail server that you want to use to send e-mails generated by OmniVista. You can also specify the "From" address that will be used for these e-mails. Complete the fields and click on the **Apply** button. The change takes effect immediately.

- **SMTP Server** - The Host Name or IP address of the SMPT Mail Server to be used for e-mails generated by OmniVista.
- **"From" Address** - The "From" address to be used in e-mails generated by OmniVista.
- **SMTP Authentication** - Enables (On) / Disables (Off) SMTP Authentication.
- **"To" Address To Test** - Enter an e-mail address to send a test e-mail and click on the **Send Test E-Mail** button.

Note: All of the fields in the must be filled or the e-mails you define will not be sent. Mail servers usually require the "From" address to be a valid e-mail address. If it is not, the mail server is likely to discard the request.

Note: OmniVista can be configured to generate and send an e-mail upon receipt of user-specified traps. This can be configured from the Automatic Trap Responders window in the Notifications application. The "To" address for Trap Responder e-mails is specified in the Automatic Trap Responders window. The "From" address and the mail server to be used are specified as shown above.

Install Zulu CEK

The Preferences Install Zulu CEK Screen is used to install the Zulu Cryptography Extension Kit (CEK), which is required to support SNMPv3 with AES 192 and 256 authentication protocols. Click on the link to read through "Zulu Terms of Use". Click on the "Download the Zulu CEK" link and download the CEK Zip File. The Zip File contains a README.txt file, a License file, a Disclaimer file and two ".jar" files (local_policy.jar and US_export_policy.jar). Click on the **Browse** button to locate and select the two (2) ".jar" files. Click on the **Upload** button to upload and install the files.

ProActive Lifecycle Management Overview

The ProActive Lifecycle Management (PALM) Feature periodically gathers detailed information for all discovered devices on your network and periodically uploads the information to the PALM web portal. Basic inventory information is also available through a widget that can be displayed on the OmniVista 2500 NMS Dashboard for easy reference. When you install OmniVista 2500 NMS, the PALM option is selected by default on the License Agreement Screen. If you accept this default, the feature is enabled and information is gathered and sent to the web portal. After the initial upload, updated information is sent to the web portal every two (2) weeks.

Note: You must have a PALM account set up before using the feature. Contact your Business Partner or Alcatel-Lucent Enterprise Customer Support for more information. Also note that if you choose not to enable the PALM Feature at installation, you can enable it at a later time in the Preferences application.

Setting Up ProActive Lifecycle Management

After installing OmniVista you must enable PALM (if necessary), and add the ProActive Lifecycle Management Widget on the OmniVista Dashboard.

Enabling PALM

During OmniVista 2500 NMS installation, users have the option to enable PALM on the License Agreement Screen. If you enabled PALM at installation, go to Adding the ProActive Lifecycle Management Widget. Otherwise, follow the steps below to enable the feature.

1. Go to the Preferences - System Settings - ProActive Lifecycle Management Screen.
2. Click on the **ProActive Lifecycle Management** slider to enable the feature. Accept the License Agreement and click the **Apply** Button.
3. After clicking **Apply**, a "Verify Proxy Configuration" link will appear to enable you verify/update your proxy settings. If necessary, click on the link to go to the Preferences - System Settings - Proxy Screen to view/change proxy settings.

Note: In addition to the PALM option presented during OmniVista 2500 NMS installation, you will be prompted to enable PALM whenever you add/relicense an OmniVista Core License in the License application.

Adding the ProActive Lifecycle Management Widget

The PALM Widget can be displayed on the OmniVista 2500 NMS Dashboard. The widget provides basic inventory information. Follow the steps below to add the widget to the Dashboard.

1. On the OmniVista 2500 NMS Dashboard, click on the Settings icon and select **Add Widget**.

Adding the Widget

The screenshot shows the OmniVista 2500 NMS dashboard interface. At the top, there is a navigation bar with the Alcatel-Lucent logo and various menu items like Home, admin, Help, Videos, About, and Logout. Below the navigation bar, the main dashboard area is titled 'OmniVista 2500 NMS'. A 'Settings Icon' (gear) is located in the top right corner of the dashboard area. A context menu is open over this icon, showing options: '+ Add Widget' and 'Change Layout'. The dashboard itself contains several widgets: 'ProActive Lifecycle Management' (a pie chart showing 91.11% and 88.89%), 'Top N Apps Summary View' (a pie chart with segments for 43.70%, 30.38%, and 14.00%), 'Top N Clients Summary 1' (a pie chart with segments for 38.60% and 38.34%), 'Top N Ports Utilization Detail' (a bar chart), 'Licenses' (a table showing expiration dates for Alcatel-Lucent Devices, Third Party Devices, and VMs), and 'Network Availability' (a donut chart showing 25 Up, 0 Warning, and 0 Down).

2. Scroll down the list of widgets, select ProActive Lifecycle Management and click **OK**. The widget will appear on the Dashboard.


Note: The ProActive Lifecycle Management Widget is activated after the initial data upload, which occurs within two weeks of installation. Once the initial upload is complete, information is displayed in the widget.

Viewing Information

Information is displayed in the ProActive Lifecycle Management Widget. A series of pie charts provide an overview of devices on your network. For a more detailed view, you can click on the widget to go to the PALM web portal.

ProActive Lifecycle Management Widget

The ProActive Lifecycle Management Widget includes a series of pie charts that provide a quick overview devices on your network. The example below shows the Operating System Release Screen. Hover over a section to view basic information (e.g., number of network devices supported/not supported). Click on the arrows (>>) to scroll through the different screens. Click on a pie chart to go to the PALM web portal for a more detailed view. The information in each pie chart is defined below.



- **Operating System Release** – Provides an overview of network devices running AOS software by displaying the percentage and number of network devices running supported and unsupported AOS software. "Not Supported" indicates that the software version running on a device is no longer supported by Alcatel-Lucent Enterprise.
- **Hardware Lifecycle** - Provides a "lifecycle" overview of network devices. "Not Supported" indicates a device has passed its End of Life date.
- **Warranty Status** - Provides an overview of the warranty status of network devices. "Not Supported" indicates that a device has past the warranty end date.
- **Support Service** – Provides an overview of the Support Agreement status of network devices. "Not Supported" indicates that a device no longer has a valid Support Service (maintenance) agreement.

PALM Web Portal

The ProActive Lifecycle Management web portal provides detailed information for your network devices. Click on any of the pie charts in the ProActive Lifecycle Management Widget to go to the web portal (login is required the first time you access the portal). The initial screen displays the same pie charts displayed on the OmniVista Dashboard. Click on any of the pie charts and scroll down to the bottom of the screen to view

detailed information.

Manually Uploading Information

After the initial data upload, OmniVista 2500 NMS sends updated information to the PALM web portal every two weeks. However, you can manually initiate a data upload at any time on the ProActive Lifecycle Management Screen in the Preferences Application. Go to the Preferences - System Settings - ProActive Lifecycle Management Screen and click on the **Upload Now** button.

Detailed Information

The table below provides a list of the specific information gathered and uploaded by the PALM Feature.

Information Category	Description
OV Server	Version of OmniVista
	MAC Address of OmniVista
	OmniVista ID
	OmniVista License Key
Device	Device Name
	Device License Key
	Device Model
	Device MAC Address
	Device Modules
	Device Last Known Update Time
	Device Chassis
	Chassis ID
	Chassis Serial Number
	Chassis Administrative Status
	Chassis Operational Status
Modules	Module MAC Address
	Module Serial Number
	Module Name
	Module Description
	Module Uboot Version
	Module Index
	Module Hardware Version
	Module Model
	Module Firmware Version
Module Software Version	

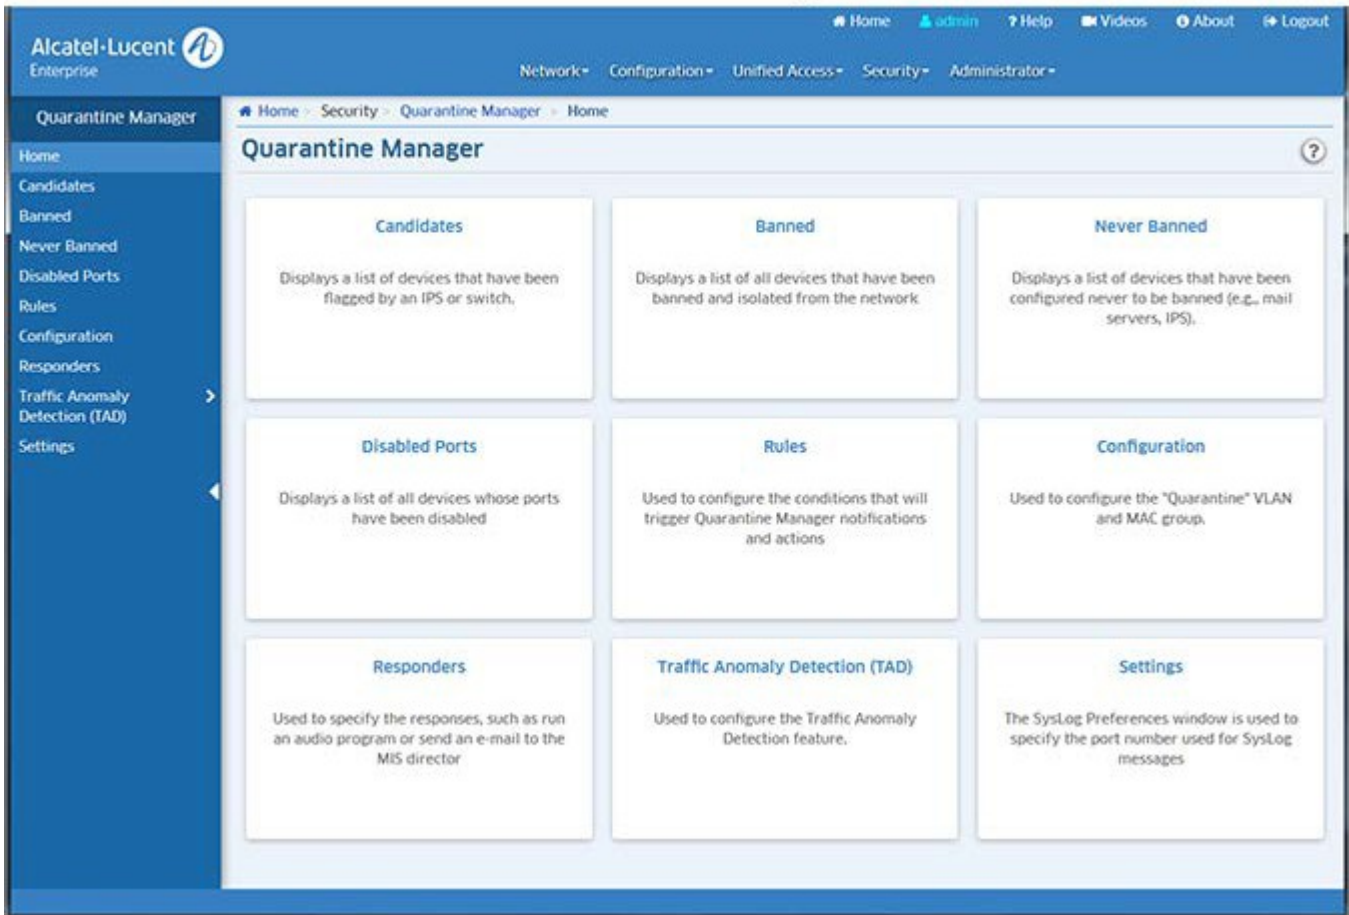
19.0 Quarantine Manager

The Quarantine Manager application enables the network administrator to quarantine devices to protect the network from attacks. The application works with an external Intrusion Prevention System (IPS), such as Fortinet, or a network device, such as an Alcatel-Lucent Enterprise AOS switch, which sends either a Syslog message or SNMP trap to Quarantine Manager containing the IP or MAC address of the offending device. (If an IP address is received, Quarantine Manager uses its Locator function to determine the device's MAC address.) These messages trigger Quarantine Manager Rules. Depending on the rule that is written for the event, the device can be immediately quarantined or placed in a Candidate List that can be reviewed by the Network Administrator for further action.

The application also includes the optional Quarantine Manager Remediation (QMR) feature. QMR is a switch-based application that interacts with Quarantine Manager to restrict network access of quarantined clients and provide a remediation path for these clients to regain their network access.

Note: Quarantine Manager cannot quarantine any devices on the EMP subnet because the EMP port has no mobility feature.

Quarantine Manager



Quarantine Manager is configured using the following screens, which can be accessed from the Quarantine Manager Home Page or by clicking on the links on the left side of any Quarantine Manager screen.

- **Candidates** - Displays a list of devices that have been flagged by an IPS or switch. The Network Administrator can release a device from the list, ban a device, or configure a device to never be banned.
- **Banned** - Displays a list of all devices that have been banned and isolated from the network.

- **Never Banned** - Displays a list of devices that have been configured never to be banned (e.g., mail servers, IPS). Note that all switches discovered by OmniVista are implicitly on the Never Banned List even though they are not displayed.
- **Disabled Ports** - Displays a list of all devices whose ports have been disabled.
- **Rules** - Used to configure the conditions that will trigger Quarantine Manager notifications and actions.
- **Configuration** - Used to configure the "Quarantine" VLAN and MAC group, as well as the action that will be taken for the event. It is also used to configure the optional Remediation Server.
- **Responders** - Used to specify the responses, such as run an audio program or send an e-mail to the MIS director, based on the conditions given. Quarantine Manager Logs
- **Traffic Anomaly Detection (TAD)** - Used to configure the TAD feature. TAD is a network monitoring feature that detects anomalies in the network traffic by monitoring the difference in the rate of ingress and egress packets on a port, matching a specific traffic pattern.
- **Settings** - Used to specify SysLog settings for Quarantine Manager. The Audit application can be used to access the Quarantine log and Syslog. The logs contain detailed information about Quarantine Manager and Syslog events.

The following sections detail Quarantine Manager requirements and basic Quarantine Manager configuration.

Quarantine Manager Requirements

The following sections detail hardware/software and basic configuration requirements for Quarantine Manager.

Hardware/Software Requirements

OmniVista Hardware/Software

Quarantine Manager is supported on AOS6250, 6350, 6400, 6450, 6850, 6855, 6860, 9000, and 10K Switches, as well as Aruba. Fortinet software version 2.3 is supported.

External Notification Device

An external device must be set up to send notifications to the Quarantine Manager application. The application works with an external Intrusion Prevention System (IPS), such as Fortinet, or a network device, such as an Alcatel-Lucent Enterprise AOS switch, which sends either a Syslog message or SNMP trap to Quarantine Manager containing the IP or MAC address of the offending device.

For example, a Fortinet IPS device must be set up to send Syslog messages to Quarantine Manager. This set up includes specifying the IP address of the OmniVista server and the port address for the OmniVista Syslog daemon (preset default is 514); and specifying what events received by the IPS will generate a Syslog Message. The message (either Syslog message or trap) must include the IP or MAC address of the offending device.

Remediation Server (Optional)

You can set up a Remediation Server that will work with Quarantine Manager to notify the user when a device is placed into the Banned List. The feature can also be configured to utilize programs/patches to debug the device and restore network access.

Quarantine Manager Remediation (QMR) is a switch-based application that interacts with Quarantine Manager to restrict network access of quarantined clients and provide a remediation path for these clients to regain their network access. A Network Administrator can set up a Remediation Server that will work with Quarantine Manager to notify the user when a device is placed into the Banned List. The feature can also be configured to utilize programs/patches to debug the device and restore network access.

When Quarantine Manager quarantines a client, the client MAC address is added to the MAC address group on the LDAP server. QMR pulls the MAC addresses from this group to populate the "Quarantined" MAC address group on the switch. At this point, network access for these clients is restricted to communication with the designated Remediation Server, and exception subnet if configured (and essential protocols such as ARP, DHCP, and DNS), until the client's quarantined status is corrected. When a client has corrected its quarantined state, Quarantine Manager updates the MAC address group on the LDAP server to remove the MAC address of the client. QMR will then restore network access to that client the next time QMR checks the LDAP MAC address group. For more information, [click here](#).

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.

Configuration Requirements

Quarantine Manager is configured using the Configuration and Rules Screens as described below. Detailed instructions for each screen are provided in the on-line help for the screen.

Quarantine VLAN and MAC Group

A basic "Quarantine" VLAN is pre-configured on OmniVista. You customize this "basic" Quarantine VLAN using the Configuration Screen. You must also configure a "Quarantine MAC Group" using the Groups application. After configuring the Quarantine VLAN and MAC Group, you apply the configuration to devices you want to monitor using Quarantine Manager.

Quarantine Manager Rules

Quarantine Manager Rules are configured for dealing with Syslog events and SNMP traps. The easiest way to use Quarantine Manager is to enable one of the Built-in Rules. The Rules determine which events from an external IPS or switch are propagated through the network. For example, when the IPS notices an attack, it generates a Syslog event. After receiving a Syslog message, Quarantine Manager uses the rules to determine what device generated the event and whether or not the offending device is immediately quarantined (Banned) or placed on the Candidate List to be reviewed by the Network Administrator. The way in which a device is quarantined depends on the action that is configured for the rule.

If a device is placed in the Candidate list, all traffic to the suspect device continues. The Network Administrator reviews each event in the Candidate List and decides what action to take. If a device is placed in the Banned list, it is quarantined until it is manually removed by the Network Administrator.

Note: There are a number of important devices in a network that a Network Administrator will never want to be quarantined. Use the Never Ban List to ensure that important devices are never quarantined.

Creating Quarantine Subnets (Optional)

If a device is banned either by the Network Administrator or Quarantine Manager, the ban is applied to all devices in the network. However, you can segment your network by creating a logical "Quarantine" network.

This will limit Quarantine Manager actions to only those switches in the "Quarantine" subnetwork(s). To create "Quarantine" subnet(s) create a map in the Topology application "Quarantine". You then create Quarantine subnets under the Quarantine network.

Candidates

The Quarantine Manager Candidates Screen displays all devices that have been placed in the Candidates Quarantine List. When an external Intrusion Prevention System (IPS), such as Fortinet, detects a possible attack on the network, it generates either a Syslog Event or an SNMP Trap. A Quarantine Manager rule can be configured (Configuration Screen) to trigger an action based on these events. The action will either immediately quarantine the offending device, or place the device on the Candidates Quarantine List. If the device is placed on the Candidates List, traffic to and from that device will continue until the Network Administrator decides what action should take place. The Network Administrator can:

- **Release the Device from the Candidates List** - To remove a device from the Candidates list, select the device and click the **Release** button. The device is removed from the list. A device may return to the list if another event triggers a configured quarantine rule.
- **Ban the Device** - To ban a device from the network, select the device and click the **Ban** button. The device is removed from the network and placed in the Banned Quarantine List.
- **Place the Device on the List of Devices to Never be Banned** - To place a device in the Never Banned list, select the device and click the **Never Ban** button. The device is placed in the Never Banned list. An event will never trigger a quarantine rule for a device in the Never Banned Quarantine List.

Candidates Quarantine List

- **MAC Address** - The device's MAC address. Quarantine Manager Rules extract the IP address of the device. Quarantine Manager then uses the OmniVista Locator Function to determine the MAC address.
- **IP Address** - The device's IP address. All Quarantine Manager Rules must extract the IP address from the Syslog Message or SNMP Trap. If the IPS sends a MAC address, the IP address will have a value of 0.0.0.0.
- **Timestamp** - The date and time the event occurred.
- **Reason** - The reason the event triggered a Quarantine Manager rule. For all Fortinet-generated events, select the event in the table and right-click to access a detailed description of the event. Click [here](#) for more information.
- **Incident Count** - The number of times an anomaly has been seen for the candidate device.

Fortinet Web Site

You can access the Fortinet web site for a detailed description of any Fortinet event. To access the description:

1. Click on the event in the Candidates Quarantine List to highlight it; then right-click on the event. The Reason window will appear.
2. Click on the Fortinet web address button at the bottom of the Reason window. A Fortinet In-Depth Analysis page will appear describing the event in detail and providing any recommended actions.

Banned

The Quarantine Manager Banned Screen displays a list of all devices that have been quarantined, either by a Quarantine Manager rule, or by the Network Administrator. When a device is placed in the Banned Quarantine List, it is quarantined from the rest of the network. Devices can automatically be added to the Banned List based on a Quarantine Manager rule or manually placed in the list by the Network Administrator. Once a device is placed in the Banned List, it remains quarantined until the Network Administrator manually releases it.


A Network Administrator can add a device to the list, edit a device on the list, release a device from the list, retry adding a device to the list, or re-poll the network for banned devices.

Note: DHCP requests from a banned device are sent to the Quarantine VLAN. The Network Administrator can direct banned traffic from the Quarantined VLAN to a Remediation Server that will provide the user with information explaining why their device was banned and what steps to take to connect to the network.


Note: Quarantine Manager can ban devices connected to an OmniAccess WLAN device using the device's "Blacklist" feature. However, the 'enable' password of the device must be entered in the Secondary Password field for the device using the "Discovery - Edit Device" Operation in the Topology application.

Note: Quarantine Manager uses a "Fast Re-Cache" mechanism. With this mechanism, the switch will look through LDAP only for the existence of quarantine MAC groups. The contents of the MAC group are added to the quarantine settings without flushing any other policies. This feature is only available on the 6400, 6850, 6855 and 9000 Series Switches running 6.3.1.R01 or later.

Adding a Device to the Banned List

In addition to automatically quarantining devices based on a Quarantine Rule, you can also manually quarantine a specific device by adding it to the Banned List. Click on the Create icon . Select the IP Address or MAC Address option, enter the IP address (or Host Name) or MAC address. Enter an optional explanation in the Reason field and click on the **Create** button. The device will appear in the list.

Editing a Device on the Banned List

Certain information about a banned device may not be picked up by a QM Rule. A Network Administrator is allowed to edit the IP Address and Reason for an entry in the Banned List to make it more closely match what the Network Administrator knows to be the best information about a banned device. Select the device in the list and click on the Edit icon . Edit the field(s) and click on the **Apply** button.

Releasing a Device from the Banned List

To release a device from the Banned List, select the device(s) in the list and click on the **Release** button. Click **OK** at the confirmation prompt.

Retry

To retry a failed operation (e.g., release device from the Banned List), select the device in the Banned List and click on the **Retry** button.

Re-Polling for Banned Devices

Click on the **Redo Ban** button to poll the network for banned switches. This is useful if you have banned switches without first creating a Quarantine VLAN or MAC Group.

Banned Quarantine List


- **MAC Address** - The device's MAC address. Quarantine Manager Rules extract the IP address of the device. Quarantine Manager then uses the OmniVista Locator function to determine the MAC Address.
- **IP Address** - The device's IP address or the host name. All Quarantine Manager rules must extract the IP address from the Syslog Message or SNMP Trap.
- **Timestamp** - The date and time the event occurred.
- **State** - The state of the banning action: Scheduled to be Banned (Ban is in process) Completed (Ban is complete) Partially Banned (Ban not completed for all devices) Scheduled to Be Released (Release is in process) Partially Released (Release not completed for all devices).
- **Reason** - The reason the event triggered a Quarantine Manager rule. For all Fortinet-generated events, select the event in the table and right-click to access a detailed description of the event. [Click here for more information.](#)
- **Partial Results** - The devices where the ban has either succeeded, or for which the user has not configured/enabled Quarantine Manager. You can click on an entry to display a detailed view listing the switches that are included in the quarantine of the banned device.
- **VLAN Name** - The user-configured name for the Quarantine VLAN.
- **MAC Group Name** - The user-configured name for the Quarantine MAC Group.

Never Banned

The Quarantine Manager Never Banned Screen displays a list of devices that have been specified by the Network Administrator never to be banned (e.g., mail servers, IPS). The screen is used to add a device to the list, edit a device description on list, and delete a device from the list. A device placed on the Never Banned list can never be banned, either manually or automatically by Quarantine Manager. Important network servers should be placed in the Never Banned Quarantine List.


Note: The OmniVista Server and all switches discovered by OmniVista are implicitly placed in the Never Banned list. Even though these devices do not appear in the list, they cannot be banned.

Adding a Device to the Never Banned List


Click on the Create icon  and complete the fields as described below. When you are finished, click on the **Create** button.

- **Address** - Select the applicable tab and enter the device's IP Address (or host name) or MAC Address. Note that you can enter the host name only if the IP Address radio button is selected. If you ban a device by the MAC address, the IP address will display a value of 0.0.0.0. If you ban a device by its IP address, Quarantine Manager will use its Locator function to determine the MAC address.
- **Reason** - Enter a reason for placing the device on the Never Banned List.

Editing a Device on the Never Banned List

You can edit the Reason field for a device in the Never Banned List. Click on the Edit icon . Edit the field(s) as described above and click on the **Apply** button.

Deleting a Device from the Never Banned List

Select a device in the Never Banned List and click on the Delete icon . Click **OK** at the Confirmation prompt.

Never Banned Quarantine List

The Never Banned Quarantine List provides basic information for all devices on the list. Click on a device to display more detailed information.

- **MAC Address** - The device's MAC address.
- **IP Address** - The device's IP address or host name. If an Intrusion Prevention System (IPS) sends a MAC address, the IP address will have a value of 0.0.0.0.
- **Timestamp** - The date and time the device was placed on the Never Banned list.
- **Reason** - The reason the device is in the Never Banned list.

Disabled Ports

The Quarantine Manager Disable Ports Screen displays a list of all devices whose ports have been disabled, either by a Quarantine Manager Rule or by the Network Administrator. The screen is used to release a device from the list, edit a device description on list, and retry a failed port operation. When a port is disabled, an entry appears in the Disabled Ports List table. If you attempt to ban multiple MAC addresses for the same switch's slot/port, multiple entries will appear in the table.


Note: If you disable a port that was already disabled, there will be two entries in the table. The first entry will contain the MAC address of the offending end station. The second entry will contain a null (possibly 000000:000000) MAC address. The reason for this second entry is that when you use the Banned Screen to release a MAC address, the port will not be re-enabled. The Network Administrator will have to manually re-enable the port by releasing the port from the Disabled Ports List.

Note: When you release an entry from the Disabled Port List, the item will be removed. If it is the last item with the specified IP address and slot/port combination, then that port will be enabled. That is, the port will not be enabled until every device that caused it to be banned has been released.

Release a Device from the Disabled Port List

Select a device(s) in the list and click on the **Release** button. Click **OK** at the confirmation prompt.

Edit a Device in the Disabled Port List

You can edit the Reason and the Timestamp fields for a device in the Disabled Ports List. Select a device in the list and click on the Edit icon . Edit the field(s) and click on the **Apply** button.

Retry a Port Operation

To retry a failed operation during the enabling/disabling of a port, select the device in the Disabled Ports List and click on the **Retry** button.

Disabled Port List

- **Switch Address** - The device's IP address. Quarantine Manager Rules extract the IP address from Syslog Message or SNMP trap.
- **Port** - The disabled slot/port number.
- **Timestamp** - The date and time the event occurred.
- **MAC Address** - The device's MAC address. Quarantine Manager Rules extract the IP address of the device. Quarantine Manager then uses the OmniVista Locator function to determine the MAC address.
- **State** - The state of the disabling or enabling action:
 - Completed (Disabling is complete)
 - Failed (Disabling/enabling of a port failed)
- **Reason** - The reason a port was disabled.

Rules

The Quarantine Manager Rules Screen displays all configured Quarantine Manager Rules and is used to configure the conditions that will trigger Quarantine Manager notifications and actions. Quarantine Manager Rules determine which Syslog events or SNMP traps cause a device to be placed in the Candidates List or Banned List, or released. You can create, edit, delete, enable/disable, and import rules.

Quarantine Manager Rule Overview

Quarantine Manager Rules determine the conditions that will trigger Quarantine Manager notifications and actions. A rule contains:

- A trigger expression that specifies the event or trap that will trigger an action
- An extraction expression that is used to extract the source address from the event or trap
- An action to be taken when the event or trap is received (device is placed in the Candidates List or Banned List, or released).

Note: Banned rules have precedence over Candidate rules. If an event matches more than one rule, Quarantine Manager will match the first rule that places a device on the Banned list. If there is no rule that places the item on the Banned list, Quarantine Manager will match the first rule that places the device on the Candidate list.

Rule Types

There are two types of rules: Built-In Rules and Custom Rules. The Built-In Rules cannot be deleted (although they can be modified or disabled). Custom Rules are rules that the Network Administrator creates.

Built In Rules

There are thirteen (13) Built-In Rules that come with Quarantine Manager. Built-In Rules are initially configured in the Disabled state. You must edit these rules to change the "Enabled" status to "True" to enable these rules. The default action configured for all of the Built-In Rules is to send the device to the Candidates list for review by the Network Administrator. Although the rules are pre-configured, the Network Administrator can modify the them. The Built-In Rules are:

- **Alcatel DOS Trap Rule** - Triggers an action based on an AOS DOS trap (AlaDosTrap). The rule triggers an action in response to a Teardrop, Ping of Death, or Port Scan attack. You can use Regular Expressions to create rules for additional AOS DOS traps.

- **Brick** - Triggers an action on a Brick Anomaly Event.
- **Fortinet Anomaly** - Triggers an action on a Fortinet Attack Anomaly Event. Ignores Anomaly attacks configured to "Pass" on Fortigate.
- **Fortinet Signature** - Triggers an action on a Fortinet Syslog Signature event. Ignores Signature attacks configured to "Pass" on Fortigate.
- **Fortinet Virus** - Triggers an action on a Fortinet Virus Detection event. Only triggers on sub-type "infected".
- **HTTP Server DOS Attack Trap** - Triggers an action when a "Denial of Service" Trap is received from an HTTP Server.
- **OA SafeGuard Malware Cleared** - Triggers and action when SafeGuard clears malware.
- **OA SafeGuard Malware Detected** - Triggers and action when SafeGuard detects malware.
- **OA WLAN Containment on AP** - Containment has been enabled for a suspected rogue AP because the confidence level for that AP equals or exceeds the configured value for that setting.
- **OA WLAN Potential Rogue AP** - An AP has been detected with conditions that may cause it to be classified as a rogue or suspected rogue.
- **OA WLAN: Rogue AP Active** - Triggers an action when the switch classifies an Access Point as a "Rogue AP."
- **OA WLAN: Rogue AP Detected** - Triggers an action when the Access Point detects an active "Rogue AP."
- **OA WLAN: Station w/ Rogue AP** - Triggers an action when the Access Point detects traffic from a client through a "Rogue AP."

Fortinet Anomaly and Signature attack events include a "status=" attribute that can be "clear_session", "pass_session", "dropped", "reset", or "detected". When Fortigate is configured to allow a particular attack (using the GUI to set its action to "Pass"), a Syslog event is still sent out for that attack, but its status is "detected"; meaning it is detected but not acted upon. Our built-in triggers are therefore designed to act on any value of "status=" EXCEPT for "detected". This means you can use the Fortigate control panel to selectively enable or disable attack actions and Quarantine Manager will behave consistently, without the need to change any of these triggers.

The Canned rules in Quarantine Manager for Fortigate have been modified as such, in both the anomaly and signature rules.

- log_id=0421073001.*status=[^p].[^t]
- log_id=0420070000.*status=[^p].[^t]

The [^p] was added to exclude any Syslog message starting with a "p" character, as well as a "t" character. This prevents quarantine for both "detected" and "pass_session" status.

The "pass_session" state was previously unknown.

Note: Built In Rules cannot be deleted, however, they can be edited and enabled/disabled.

Note: The Audit application can be used to access the Quarantine log and Syslog. The logs contain detailed information about Quarantine Manager and Syslog events.

Custom Rules

The Network Administrator can create Custom Rules using Regular Expressions to configure the trigger event and extraction expression. The rules can be based on an Intrusion Prevention System (IPS) event or an AOS SNMP trap notification. Custom Rules can be created, edited, deleted, enabled/disabled, and imported.

Note: You must be careful when creating a rule since a mis-configured rule could cause an important service to be inadvertently banned.

Regular Expressions Overview

Trigger Expressions

As stated earlier, a Trigger Expression is a regular Java expression that is used to determine if a Syslog message or SNMP Trap should trigger a quarantine action. If a Syslog Message or SNMP Trap matches this regular expression, the action is performed. For example, let's say that we are interested in a Fortinet Syslog Event that looks something like:

```
Fortinet Anomaly 03-08-200 14:09:34 device_id=FG36002805033253 log_id=0421073001 type=ips
subtype=anomaly pri=critical attack_id=102039582 src=90.0.0.10 dst=10.10.10.100 src_port=2370
dst_port=139 src_int=internal dst_int=external status=dropped proto=6 service=139/tcp msg="netbios:
SMB.NTLMSSP.Attempt.B
```

Many Syslog messages appear similar. However, each message may have a different date, device ID, source and destination address, etc. What is unique about each Syslog message is the log_id value. If you are interested in all Fortinet Syslog messages with a log_id of 0421073001, then the regular expression is easy - you can simply search for any message that contains the String log_id=0421073001. In the Trigger Expression Field you would type the value

```
log_id=0421073001.
```

Extraction Expressions

As stated earlier, an Extraction Expression is a regular Java expression that specifies the source address of the offending device. Once a Syslog message matches a Trigger Expression, Quarantine Manager must extract the source address of the suspect end station from the message. In the Fortinet example above, the source address is preceded by the string "src=" and then an IP address. An IP address consists of 4 sets of numbers separated by the "." character. Each set is 1 to 3 characters in length and the numbers are decimal (0-9) digits. One way to express this is with the regular expression

```
src=([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})
```

Because we are only interested in the IP address and not the characters "src=", place () around the IP address to indicate which part you want to capture. The [0-9] means any single character from 0-9, the {1,3} means that you are looking for a set of 1 to 3 numbers. The \. says that you are looking for a "." character.

The backslash is an escape character that says take the "." literally (normally "." is a special character that means any character).

Often there are a number of regular expressions that you can use to achieve the same results. In the Fortinet example above, there is the string "src=" followed by the IP address, followed by either a space or a tab character. The regular expression for getting the IP address could be

```
src=( [^ ]*)
```

The characters between the [] are ^, a space character and a tab character. This expressions says: the string "src=" followed by a sequence of characters that are not spaces or tabs. This expression works well, but it can be difficult to read because the space and tab character are not "visible".

Another way to extract the IP address would be src=([0-9.]*)

This says, the string "src=" followed by a sequence of characters that contain only numbers and the "." (a "." inside of [] does not need to be escaped).

Note: If an extraction expression is not working, check the server.txt file to troubleshoot the problem.

Useful Operators for Quarantine Manager Rules

^ the "not" operator

[start of character list

] end of character list

(beginning of an expression

) end of an expression

[.] useful for escaping characters used as operators in regular expressions

Basic Regular Expressions

. Matches any single character

[...] Matches any one of the characters enclosed between the brackets. If the first character is a circumflex (^), then it matches any one character Not enclosed between the brackets. A hyphen (-) is used to indicate a range of characters.

\ Escape the special character that follows.

* Matches any number (including none) of the single character that immediately precedes it.

+ Matches one or more occurrences of the preceding regular expression.

? Matches zero or one occurrences of the preceding regular

expression. For example:

[abc] Matches either an 'a', 'b' or

'c' [a-z] Matches all lower case

letters [a-zA-Z] Matches all letters

[0-9a-fA-F] Matches all hex digits

[^0-9] Matches any character that is not a digit

AOS DOS Trap Configuration

In addition to the Built-In Rule for AOS DOS Traps (0, 2, 6), you can configure a Rule for other AOS DOS traps. For example, the built-in rule for AOS DOS Traps is:

```
TrapName=alaDoSTrap.*alaDoSType=[0|2|6]
```

This triggers a response on AOS DOS Trap type 0, 2, and 6.

To trigger on types 0, 2, 6, 9, 10, 11, 12 and 13, you would enter:

```
TrapName=alaDoSTrap.*alaDoSType=(|[0|2|6|9|10|11|12|13])
```

The () form a group

Inside the group are basically two parts separated by a vertical bar | which means either or so we have

(A | B)

The first part (A)

is [0|2|6|9]

The square brackets [] mean match a single character from the list. The vertical bar | means or. So this expressions says either a 0 or 2 or 6 or 9.

The second part (B)

is 1[0123]

which means a 1 followed by either a 0 or a 1 or a 2 or a

3 That expression could have been written as

1[0-3]

The dash - is a special character which is used to express a range.

Note that we could **not** have written the expression as alaDoSType=[0-13] to match all 13 types.

This expression says a 0 through 1 or a 3. so it would match:

alaDoSType

=0


alaDoSType

=1

alaDoSType

=3

Creating a Quarantine Manager Rule


Click on the Create icon  and complete the fields as described below. When you are done, click on the **Create** button.

- **Name** - The user-defined name for the rule.
- **Description** - The user-defined description for the rule.
- **Trigger Expression** - A regular Java expression that is used to determine if a Syslog message or SNMP trap should trigger a quarantine action. If a Syslog message or SNMP Trap matches this regular expression, the action is performed. The regular expressions used by OmniVista are very similar to those used by programs such as PERL and AWK. [Click here for more information on Regular Expressions.](#)
- **Extraction Expression** - A regular Java expression that specifies the source address of the suspect device. Use the () expression to capture the source IP or MAC address. (Quarantine Manager also supports the hex form of IP addresses.) Once Quarantine Manager receives a Syslog message or SNMP trap that matches a Trigger Expression, it must extract from it the source address of the suspect end station. [Click here for more information on Regular Expressions.](#)
- **Action** - The action to be taken when the rule is triggered:
 - **Candidate List** - The device is added to the Candidates list. The device can still send and receive traffic. The Network Administrator reviews the list and determines what action to take (e.g., remove the device from the list, ban the device)
 - **Quarantine** - The device is moved to the Quarantined VLAN and/or MAC Group, and added to the Banned list. While on the Banned list, the device cannot send or receive traffic. The device remains on the list until it is manually removed by the Network Administrator.
 - **Release** - The device is released from the Quarantined VLAN and/or MAC Group. This can be


used to allow an external system (e.g., Trouble Ticket System) to send a syslog message or trap to OmniVista to automatically release a quarantine without having to access OmniVista. Note that the Quarantine VLAN or MAC group must be properly set up for traffic to be quarantined. If you do not first configure a Quarantine VLAN or MAC group, even a device on the Banned list could still pass traffic. The Quarantine VLAN or MAC group is configured in the Configuration Screen.

- **Event Type** - The type of triggering event (Syslog or Trap).
- **Enabled** - Administrative state of the rule:
 - **On** - The rule is enabled.
 - **Off** - The rule is disabled.

Editing a Quarantine Manager Rule

Select a rule in the list and click on the Edit icon . Edit the field(s) as described above and click on the **Apply** button. Note that you cannot edit a rule name.

Deleting a Quarantine Manager Rule

Select the rule(s) in the list that you want to delete and click on the Delete icon . Click **OK** at the Confirmation Prompt.

Enabling/Disabling a Quarantine Manager Rule

To enable/disable a rule(s), select the rule(s) in the list and click on the **Enable/Disable Rules** button.

Importing a Quarantine Manager Rule

You can import a new rule from Alcatel-Lucent Enterprise without having to update the Quarantine Manager code. New rules are sent by Alcatel-Lucent Enterprise as .xml files.

1. Save the .xml file on your machine.
2. Click on the **Import** button.
3. Browse to the .xml file that you saved, select the rule and click **OK**. The Import window will close and the new rule will appear in the Rules table. A sample of an imported .xml file is shown below.

Note: Imported rules are initially configured in the Disabled state. You must change the "Enabled" status to "True" to enable the rules.

.xml import file sample:

```
<?xml version="1.0"?>
<!DOCTYPE quarantineRules[
<!ELEMENT quarantineRules (quarantineRule)*>
<!ELEMENT quarantineRule (desc, trigger, extract)>
<!ELEMENT desc (#PCDATA)>
<!ELEMENT trigger (#PCDATA)>
<!ELEMENT extract (#PCDATA)>
<!ATTLIST quarantineRule name CDATA
#REQUIRED eventType (syslog | snmptrap)
#REQUIRED
type (standard |custom)
```

```

"standard" enabled (true | false)
"false"
action (ban | candidate | release) "candidate"
>
]>

<quarantineRules>
<!-- Fortinet Syslog IDS Signature Event -->

<quarantineRule name="test Signature" eventType="syslog">
<desc>test IDS Signature</desc>
<trigger>log_id=0420070000.*status=[^p].[^t]</trigger>
<extract>src=([0-9.]*)</extract>
</quarantineRule>

</quarantineRules>

```

Quarantine Manager Rule List

The Rule List displays information about all Quarantine Manager Rules stored in OmniVista.

- **Name** - The user-defined name for the rule.
- **Description** - The user-defined description for the rule.
- **Trigger Expression** - A regular Java expression that is used to determine if a Syslog message or SNMP trap should trigger a quarantine action. If a Syslog message or SNMP Trap matches this regular expression, the action is performed. The regular expressions used by OmniVista are very similar to those used by programs such as PERL and AWK. Click here for more information on Regular Expressions.
- **Extraction Expression** - A regular Java expression that specifies the source address of the suspect device. Use the () expression to capture the source IP or MAC address. (Quarantine Manager also supports the hex form of IP addresses.) Once Quarantine Manager receives a Syslog message or SNMP trap that matches a Trigger Expression, it must extract from it the source address of the suspect end station. Click here for more information on Regular Expressions.
- **Action** - The action to be taken when the rule is triggered:
 - **Candidate List** - The device is added to the Candidates list. The device can still send and receive traffic. The Network Administrator reviews the list and determines what action to take (e.g., remove the device from the list, ban the device)
 - **Quarantine** - The device is moved to the Quarantined VLAN and/or MAC Group, and added to the Banned list. While on the Banned list, the device cannot send or receive traffic. The device remains on the list until it is manually removed by the Network Administrator.
 - **Release** - The device is released from the Quarantined VLAN and/or MAC Group. This can be used to allow an external system (e.g., Trouble Ticket System) to send a syslog message or trap to OmniVista to automatically release a quarantine without having to access OmniVista. Note that the Quarantine VLAN or MAC group must be properly set up for traffic to be quarantined. If you do not first configure a Quarantine VLAN or MAC group, even a device on the Banned list could still pass traffic. The Quarantine VLAN or MAC group is configured in the Configuration Screen.
- **Event Type** - The type of triggering event (Syslog or Trap).
- **Enabled** - Administrative state of the rule:
 - **On** - The rule is enabled.
 - **Off** - The rule is disabled.

Configuration

The Quarantine Manager Configuration Screen is used to configure Quarantine Manager, including the Quarantine VLAN, Quarantine MAC Group, and the optional Quarantine Manager Remediation (QMR) feature. By default, the name of the Quarantine VLAN is "Quarantined". A basic "Quarantine" VLAN is pre-configured on OmniVista ("Quarantined"). You customize this "basic" Quarantine VLAN using the Configuration Screen.

When a Quarantine Rule extracts an IP address from a device, OmniVista uses the Locator function to determine the MAC address of the device. The device is then automatically added to the Quarantined MAC group. When devices are banned, either through a Quarantine Manager rule or by the Network Administrator, they are added to the Quarantined VLAN and/or Quarantined MAC group. These devices no longer route traffic to any other devices in the network (although you can create a logical "Quarantine" subnet to limit Quarantine Manager actions to a specific set of switches on the network.). The devices remain in the Banned list until removed by the Network Administrator.

Note: Quarantine Manager has the ability to ban devices connected to an OmniAccess WLAN device using the device's "Blacklist" feature.

Quarantined Manager Remediation (QMR)

Quarantine Manager Remediation (QMR) is a switch-based application that interacts with Quarantine Manager to restrict network access of quarantined clients and provide a remediation path for these clients to regain their network access. A Network Administrator can set up a Remediation Server that will work with Quarantine Manager to notify the user when a device is placed into the Banned List. The feature can also be configured to utilize programs/patches to debug the device and restore network access.

When Quarantine Manager quarantines a client, the client MAC address is added to the MAC address group on the LDAP server. QMR pulls the MAC addresses from this group to populate the "Quarantined" MAC address group on the switch. At this point, network access for these clients is restricted to communication with the designated Remediation Server, and exception subnet if configured (and essential protocols such as ARP, DHCP, and DNS), until the client's quarantined status is corrected.

When a client has corrected its quarantined state, Quarantine Manager updates the MAC address group on the LDAP server to remove the MAC address of the client. QMR will then restore network access to that client the next time QMR checks the LDAP MAC address group.

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.

Configuring Quarantine Manager

Configuring Quarantine Manager on network devices consists of the following steps:

- Setting Up Quarantine Manager
- Configuring Quarantine Manager
- Assigning Quarantine Manager to Devices.

Setting Up Quarantine Manager

As mentioned earlier, a basic "Quarantine" VLAN is pre-configured on OmniVista. You customize this "basic" Quarantine VLAN using the Configuration Screen. The initial Quarantine Manager setup consists of

the following steps:

- Creating the Quarantine MAC Group
- Creating the Quarantine MAC Group Policy

Creating the Quarantine MAC Group


As shown in the Quarantine Configuration List, the pre-configured OmniVista Quarantine VLAN is associated with the "Quarantine" MAC Group. However, this MAC Group is not yet configured. You must go to the Groups application and create the Quarantine MAC Group using the MAC Groups Screen. The name must match the MAC Group name associated with the Quarantine VLAN. ("Quarantined").



Note: If after creating the Quarantined MAC group you modify the name (either using the Groups application, CLI, or WebView), you must also modify the Quarantined MAC group name using the Configuration Screen and poll the switch for the change to take effect.

Creating the Quarantine MAC Group Policy

After creating the Quarantine MAC Group, you must create go to the PolicyView application and create a MAC Group Policy that will deny any traffic originating from the Quarantine MAC Group (PolicyView - Users & Groups - Unified Policies). On the Set Condition Screen of the Create Policy Wizard, create an L2 Source MAC Address Group condition for the "Quarantined" MAC Group; then on the Set Action Screen, set the disposition to "Drop".

Configuring Quarantine Manager

After creating and configuring the Quarantine MAC Group, select the Quarantine VLAN in the Quarantine VLAN List and click on Edit icon  and complete the fields as described below to complete the configuration. When you are finished, click on the **Apply** button.

- **VLAN Name** - The Quarantine VLAN name (Default = Quarantined).
- **MAC Group Name** - The Quarantine MAC Group Name (Default - Quarantined).
- **Remediation URL (Optional)** - The URL of the Remediation Server (e.g., <http://alaremediation.com>). If the Remediation Server is running on a port other than the default port for the browser (e.g., 8080), the port needs to be included in the Remediation URL (e.g., (<http://alaremediation.com:9090>)).
- **Remediation IP (Optional)** - The IP Address of the Remediation Server in the field. You must add the Remediation Server IP address to the Allowed Subnet List.
- **HTTP Proxy Port** - If there is a firewall/proxy configured for the network, enter the HTTP Proxy Port used by the network (e.g., 8080).
- **Default QMR Page** - Enables/Disables the default Quarantine Manager Remediation Web Page. If enabled, the page is automatically presented to the user if a Remediation Server is not configured.
- **Allow Port Disabling** - Enables/Disables the Port Disabling feature. You can enable or disable a port rather than a device. By default, this checkbox is disabled. If the checkbox is checked, it means that you want to disable the port when a Quarantine rule is matched. Please note that you must turn on port disabling for each device in addition to turning on global port disabling. Go to the Discovery application, select the device in the Inventory List and click on the Edit icon  to bring up the Edit Discovery Manager Entry Screen. In the Advanced Settings Section set "Allow Port Disabling" to "Yes". You can also edit a device by going to the Topology application, selecting the device, and clicking on the "Discovery - Edit Device" to bring up the Edit Discovery Manager Entry Screen. Note that port disabling looks for Locator Live Search information and does not look for historical information.
- **Subnets** - You can create an "Allowed Subnet List". This is a reserved QoS network group that includes the Remediation Server and any subnets to which a quarantined client is allowed access. 

Click on the Create icon and enter up to three (3) subnets. When you are finished, click on the **Create** button.

Note: You must add the Remediation Server IP address/subnet mask to the "QMR Allowed Subnets" Group, so that a quarantined client can communicate with the Remediation Server. You can optionally add additional subnets to which quarantined devices will have access.

Assigning Quarantine Manager to Network Devices


After completing the Quarantine Manager configuration, you must assign the configuration to network devices. Select the Quarantined VLAN in the Quarantine Configuration List and click on the **Apply to Devices** button. Select an option from the drop-down menu (Use Switch Picker/Use Topology) to select network devices and click on the **Assign** button.

Creating Quarantine Subnets (Optional)

If a device is banned either by the Network Administrator or Quarantine Manager, the ban is applied to all devices in the network. However, you can segment your network by creating a logical "Quarantine" network. This will limit Quarantine Manager actions to only those switches in the "Quarantine" subnetwork(s). To create "Quarantine" subnet(s) you use the Maps feature in the Topology application to create a Logical network called "Quarantine". You then create Quarantine subnets by creating subnetworks under the Quarantine network.

Configuring Quarantine Manager on OmniAccess WLAN Devices

Quarantine Manager can ban wireless devices connected to an OmniAccess WLAN device by placing them in the OmniAccess "Blacklist". If wireless device(s) is found in a Quarantine Segment, the MAC address of the Quarantined device is placed in the blacklist. Due to limitations in the current OmniAccess device's SNMP implementation, the banned device is placed on the blacklist using SSH to send CLI commands to OmniAccess. SSH must be able to login to the OmniAccess device. In addition, the OmniAccess device's 'enable' command must be executed and a secondary password is required to enter the privileges commands necessary to perform the blacklist. To enable automatic login, configure a valid user name and password for the OmniAccess device using the Discovery application.


Go to the Discovery application, select the device in the Inventory List and click on the Edit icon  to bring up the Edit Discovery Manager Entry Screen. In the General Section enter a CLI/FTP User Name and Password (if necessary) and enter a Secondary CLI/FTP User Name and Password. You can also edit a device by going to the Topology application, selecting the device, and clicking on the "Discovery - Edit Device" to bring up the Edit Discovery Manager Entry Screen.

Responders

The Quarantine Manager Responders Screen displays all configured Quarantine Manager automatic event responders, and is used to create, edit, and delete event responders. The screen is used to specify the response, such as external emails or scripts to be run (if any) that you want OmniVista to provide when quarantine actions are taken. This provides a method to integrate with trouble-ticket systems. OmniVista can make the following responses to the receipt of a specified event:

- Send an e-mail to any address you specify. You can use variables to specify the information you want to include in the e-mail. Variables exist for information, such as action, reason, MacAddress, etc.
- Execute an external program or script on the OmniVista server.

Creating a Quarantine Manager Responder

Click on the Create icon  and complete the fields as described below. When you are finished, click on the **Create** button.


- **Banned** - Select whether or not you want OmniVista to respond when a device is banned (Respond/Ignore).
- **Released** - Select whether or not you want OmniVista to respond when a device is released from the Banned List (Respond/Ignore).
- **Response Description** - An optional description for the Responder.
- **Response Action** - The response you want OmniVista to take.
 - **Send an E-Mail** - If you set the Respond Action to "Send an E-Mail", complete the fields as described below. It is important to note that all fields on the E-Mail Screen in the Preferences application must be complete, or the emails you define will not be sent.
 - **E-Mail To** - The address to which the e-mail will be sent. (The "From" address in the responder emails is determined from the entry in the Use "From" Address field in the E-mail window of the Preferences application.)
 - **E-Mail Subject** - The subject of the e-mail, which will appear in the e-mail Subject Line.
 - **E-Mail Body** - Enter the body of the e-mail in the E-mail Body field by typing in the desired text and/or the desired variables. The variables you can use are explained in the Event Variables section below. You can also accept the default email body, which is the variable \$Details\$ (explained below).
 - **Run an Application on the Server** - If you set the Respond Action to "Run an Application on the Server", complete the fields as described below.
 - **Command** - Enter the command(s).
 - **Arguments** - Enter the arguments to the command specified above, or accept the default argument, which is the variable \$MacAddress\$ (explained in the Event Variables section below).
 - **Start Directory** - The directory in which the command should be executed.
 - **Standard Input** - Enter the standard input for the command, or accept the default standard input, which is the variable \$Details\$ (explained in the Event Variables section below).

Event Variables


When sending an e-mail, you can specify the following variables in the E-Mail Body Field to automatically include the specified information:

- **\$Action\$** - The action being taken, a ban or a release.
- **\$Reason\$** - The Reason field from the QM object.
- **\$MacAddress\$** - The MAC address of the device being banned or release.
- **\$IpAddress\$** - The IP address of the device being banned or release. If the IP address is unknown it will be displayed as 0.0.0.0
- **\$VlanName\$** - The name of the VLAN that the device was banned to or released from.
- **\$MacGroupName\$** - The MAC group that the device was banned to or released from.
- **\$Details\$** - Contains a message with the Action, Mac, IP address, Vlan, and MacGroupName.

Editing a Quarantine Manager Responder

Select a Responder in the Automatic Event Responders List and click on the Edit icon . Edit the field(s) as described above and click on the **Apply** button.

Deleting a Quarantine Manager Responder

Select a Responder in the Automatic Event Responders List and click on the Delete icon . Click **OK** at the Confirmation prompt.

Automatic Event Responders List

The Automatic Event Responders List provides information on all configured Quarantine Manager Responders.

- **Banned** - The Response OmniVista takes when a device is banned (Respond/Ignore).
- **Released** -The Response OmniVista takes when a device is released from the Banned List (Respond/Ignore).
- **Response Description** - An optional description for the Responder.
- **Response Action** - The response OmniVista takes if the Responder is set to "Respond (Send an E- Mail/Run an Application).
- **E-Mail To** - The address to which the e-mail is sent, if applicable.
- **E-Mail Subject** - The subject of the e-mail, which will appear in the e-mail Subject Line, if applicable.
- **E-Mail Body** - The body of the e-mail, if applicable.
- **Command** - Application command(s), if applicable.
- **Arguments** - Arguments to the command, if applicable.
- **Start Directory** - The directory in which the command should be executed, if applicable.
- **Standard Input** - The standard input for the command, if applicable.

TAD Profile

The Quarantine Manager Traffic Anomaly Detection (TAD) Profile Screen displayed all configured TAD Profiles, and is used to create, edit, delete, and assign TAD Profiles. TAD is a network monitoring feature that detects anomalies in the network traffic by monitoring the difference in the rate of ingress and egress packets on a port, matching a specific traffic pattern. TAD monitors these packets at configured intervals, counts the packets matching certain patterns, and applies anomaly detection rules configured by the user when an anomaly exceeds user-defined thresholds (e.g., log the event, send a trap, quarantine a port).

Creating a TAD Profile

Click on the Create icon **+** and complete the fields as described below. When you are finished, click on the **Create** button. You can create up to 32 monitoring-groups. After creating the group, you then configure the anomaly you want to detect, configure a rule to execute when the anomaly is detected, and assign a port or set of ports to the TAD Group.

- **Group Name** - The name of the TAD Monitoring Group (up to 32 characters)
- **Anomaly Type** - The type of the anomaly to be enabled or disabled.
 - **All** (all) - All anomaly types are monitored.
 - **ARP Address Scan** (arpaddrscan) - Occurs when a host sends a burst of ARP requests for multiple IP addresses.
 - **ARP Flood** (arpflood) - Occurs when a host receives a burst of ARP request packets.
 - **ARP Failure** (arpfailure) - Occurs when ARP queries do not elicit ARP responses.
 - **ICMP Address Scan** (icmpaddrscan) - Occurs when multiple hosts receive ICMP echo request packets at the same time.
 - **ICMP Flood** (icmpflood) - Occurs when a host receives a burst of ICMP echo request packets.
 - **ICMP Unreachable** (icmpunreachable) - Occurs when a host receives a flood of ICMP Unreachable packets.


- **TCP Port Scan** (tcpportscan) - Occurs when a host receives a burst of TCP SYN packets for multiple TCP ports.
- **TCP Address Scan** (tcpaddrscan) - Occurs when multiple hosts receive TCP SYN packets at the same time.
- **SYN Flood** (synflood) - Occurs when a host receives a burst of TCP SYN packets on the same TCP port.
- **SYN Failure** (synfailure) - Occurs when a host receives fewer SYNACKs than SYNs it sent out.
- **SYN-ACK Scan** (synackscan) - Occurs when a host receives more SYNACKs than SYNs it sent out.
- **Fin Scan** (finscan) - Occurs when a host receives a burst of FIN packets.
- **Fin-Ack Diff** (finackdiff) - Occurs when a host sees more or fewer FINACK packets than it sent.
- **Rst Count** (rstcount) - Occurs when a host receives a flood of RST packets.
- **Count** - The number of packets that must be seen during the monitoring period to trigger anomaly detection. The valid range is 1 - 100,000. Supported anomalies and the default count for each are listed below:
 - **all** - NA
 - **arpaddrscan** - 50
 - **arpflood** - 90
 - **arpfailure** - 6
 - **icmpaddrscan** -30
 - **icmpflood** -90
 - **icmpunreachable** - 20
 - **tcpportscan** - 20
 - **tcpaddrscan** - 30
 - **synflood** - 90
 - **synfailure** - 10
 - **synackcan** - 2
 - **finscan** - 6
 - **finackdiff** - 5
 - **rstcount** - 50
- **Sensitivity** - Sensitivity of anomaly detection to deviation from the expected traffic pattern. The valid range is 1 - 100. (Default = 50)
- **Period** - The time duration to observe traffic pattern, in seconds. The valid range is 5 to 3,600. (Default = 30)
- **Anomaly State** - Enables/Disabled anomaly detection.
- **Log** - Enables/Disables logging of detected anomalies. If enabled, the anomaly information will be written to a syslog if the anomaly is detected at configured levels (Count/Sensitivity/Period). (Default = Disabled)
- **Trap** - Enables/Disables the sending of a trap when an anomaly is detected. If enabled, a trap is sent if the anomaly is detected at configured levels (Count/Sensitivity/Period). (Default = Disabled)
- **Quarantine** - Enables/Disables quarantining of the port on which an anomaly is detected. If enabled, a port is quarantined if the anomaly is detected at configured levels (Count/Sensitivity/Period). (Default = Disabled)

Assigning a TAD Profile

After configuring the monitoring group, you must assign the ports that you want to monitor to that group. TAD applies the rules to match the specific packets when a port is in a monitoring-group. These rules exist as long as the port is a member of any monitoring-group. Select a TAD Monitoring Group in the Monitoring Group List and click on the **Apply to Devices** button. Configure the fields as described below, then select the device(s) to which you want to assign the profile.

- **Group Name** - Pre-filled with the selected monitoring group.
- **Action** - Select "Assign" (default) from the drop-down menu to assign the profile. You can also remove the selected profile, or assign the selected profile and remove any others assigned to the device(s).
- **Force Port Override** - Enables/Disables port override (On/Off). Select "On" to assign all selected ports to this TAD Group (and remove them from any previously assigned groups, if applicable).
- **Select Devices** - Select an option from the drop-down menu (Use Switch Picker/Use Topology) and click on the **Add/Remove Devices** button to select devices. After selecting devices, click on a device and click on the **Add/Remove Ports** button to select ports. Repeat for additional selected devices. Click on the **Apply** button to assign the profile to devices/ports.

Editing a TAD Profile

Select a Monitoring Group in the Monitoring Group List and click on the Edit icon . Edit the field(s) as described above and click on the **Apply** button.

Deleting a TAD Profile

Select a Monitoring Group in the Monitoring Group List and click on the Delete icon . Click **OK** at the Confirmation prompt.

Monitoring Group List

The TAD Monitoring Group List provides information on all configured TAD Profiles.

- **Group Name** - The name of the TAD Monitoring Group (up to 32 characters)
- **Anomaly Type** - The type of the anomaly to be enabled or disabled.
- **Anomaly State** - Administrative state of Anomaly Detection (Enabled/Disabled).
- **Log** - Enables/Disables logging of detected anomalies. If enabled, the anomaly information will be written to a syslog if the anomaly is detected at configured levels (Count/Sensitivity/Period). (Default = Disabled)
- **Trap** - Enables/Disables the sending of a trap when an anomaly is detected. If enabled, a trap is sent if the anomaly is detected at configured levels (Count/Sensitivity/Period). (Default = Disabled)
- **Quarantine** - Enables/Disables quarantining of the port on which an anomaly is detected. If enabled, a port is quarantined if the anomaly is detected at configured levels (Count/Sensitivity/Period). (Default = Disabled)
- **Count** - The number of packets that must be seen during the monitoring period to trigger anomaly detection. The valid range is 1 - 100,000.
- **Sensitivity** - Sensitivity of anomaly detection to deviation from the expected traffic pattern. The valid range is 1 - 100. (Default = 50)
- **Period** - The time duration to observe traffic pattern, in seconds. The valid range is 5 to 3,600. (Default = 30)

TAD View

The Quarantine Manager TAD View Screen is used to view TAD configurations and anomaly statistics for specific switches in the network. Select an option from the drop-down menu (Use Switch Picker/Use Topology) and click on the **Select a Device** button to select a device to view. The following information is available.

- Monitoring Groups
- Port Ranges
- Statistics Port
- Statistics Anomaly Traffic
- Statistics Anomaly Summary

Monitoring Groups

The Monitoring Groups Table displays information for TAD Monitoring Groups configured on the selected switch.

- **Group Name** - The name of the TAD Monitoring Group.
- **Anomaly Type** - The type of the anomaly to be enabled or disabled. Supported anomalies are described below.
- **Anomaly State** - Anomaly detection administrative status (Enabled/Disabled).
- **Log** - Anomaly detection logging state (Enabled/Disabled). (Default = Disabled)
- **Trap** - Anomaly detection trap state (Enabled/Disabled). (Default = Disabled)
- **Quarantine** - Anomaly detection quarantine state (Enabled/Disabled). (Default = Disabled)
- **Count** - Configured Count parameter. This is the number of packets that must be seen during the monitoring period to trigger anomaly detection. The valid range is 1 - 100,000.
- **Sensitivity** - Configured Sensitivity parameter. This is the anomaly detection to deviation from the expected traffic pattern. The valid range is 1 - 100. (Default = 50)
- **Period** - Configured monitoring time period. The time duration to observe traffic pattern, in seconds. The valid range is 5 to 3,600. (Default = 30)

Port Ranges

The Port Ranges Table displays information on ports being monitored by a TAD Monitoring Group.

- **Group Name** - The name of the TAD Monitoring Group.
- **Start Slot/Port** - The first slot/port number in the range of ports being monitored.
- **End Slot/Port** - The last slot/port number in the range of ports being monitored.

Statistics Port

The Statistics Port Table displays anomaly pattern counts on ports belonging to TAD Monitoring Groups.

- **Slot/Port** - The slot/port being monitored.
- **Packet Type** - The type of packet being monitored.
- **Last In** - The number of incoming anomaly packets observed during the last 5 seconds.
- **Last Out** - The number of outgoing anomaly packets observed during the last 5 seconds.
- **Total In** - The total number of incoming anomaly packets observed since monitoring was enabled.
- **Total Out** - The total number of outgoing anomaly packets observed since monitoring was enabled.

Statistics Anomaly Traffic

The Statistics Anomaly Traffic Table displays the anomaly counts on ports belonging to TAD Monitoring Groups.

- **Slot/Port** - The slot/port being monitored.
- **Anomaly** - The type of anomaly.
- **Packet Type** - The type of packet being monitored.
- **Current In** - The number of incoming anomaly packets observed.

- **Current Out** - The number of outgoing anomaly packets observed.
- **Last In** - The number of incoming anomaly packets observed during the last 5 seconds.
- **Last Out** - The number of outgoing anomaly packets observed during the last 5 seconds.

Statistics Anomaly Summary

The Statistics Anomaly Summary Table displays the anomaly check summary.

- **Slot/Port** - The slot/port being monitored.
- **Anomaly** - The type of anomaly.
- **Observed** - The number of times an anomaly was observed on this port since monitoring was enabled.
- **Detected** - The number of times an anomaly was detected on this port since monitoring was enabled (the number of times the anomaly exceeded monitoring limits).

Settings

The Quarantine Manager Settings Screen is used to specify the port number used for SysLog messages. SysLog messages are used by Quarantine Manager to configure network responses. Configure a field(s) as described below and click on the **Apply** button.

SysLog Listener

- **SysLog Port Number:** The port number of the SysLog Listener (Default = 514).

SysLog Generator Target

- **SysLog IP Address:** The IP address of the device that will receive the syslog messages.
- **SysLog Port Number:** The port number on the receiving device that will receive the syslog messages.

20.0 Report

The Report Application enables you to create and schedule reports in certain OmniVista applications (e.g., Discovery, Locator, Analytics). These reports are generated and stored as PDF documents. So in addition to viewing information in real-time in OmniVista (e.g., Discovery Inventory List, Analytics Utilization Reports), you can generate PDFs of the screens. When a report is generated, it takes a current snapshot of the application information. These reports can be generated immediately or you can schedule them to be generated at regular times/intervals (e.g., Daily, Weekly). You can also configure a report to be e-mailed when it is generated. The Report Configuration Screen is used to create/configure a report. These generated reports are then displayed on the Report List Screen, where they can be downloaded and viewed as PDFs.



Note: Only users authenticated through the Local OmniVista Authentication Server can generate reports. Report generation is not supported for users authenticated through an external RADIUS Server.


Report Configuration

The Report Configuration Screen is used to create, edit, and delete Reports. These reports are PDF versions of tables and reports generated in certain OmniVista applications (e.g., Discovery, Locator, Analytics).

Basically, in addition to viewing information in real-time in OmniVista (e.g., Discovery Inventory List, Analytics Utilization Reports), you can generate PDFs of the screens. When a report is generated, it takes a current snapshot of the application information. These reports can be generated immediately or you can schedule them to be generated at regular times/intervals (e.g., Daily, Weekly). You can also configure a report to be e-mailed when it is generated. These generated reports are then displayed on the Report List Screen, where they can be downloaded and viewed as PDFs.

Creating a Report

There are two steps to creating a report. First you must configure the report in the Report Application (report name, schedule, e-mail), then you must go to an application that supports the Report Feature (e.g., Discovery, Locator, Analytics), click on the **Add to Report** button at the top of a screen, and link that report to a Report Configuration.

1. Click on the Create icon  and complete the fields as described below. After completing the fields, click on the **Create** button.

- **Report Title** - Enter a title for the report.
- **Schedule Settings**
 - **Purging Policy** - The report purging frequency. Select an option from the drop-down menu. The report will be removed from the server at the selected interval. Select "None" to never purge the report.
 - **Schedule** - The report creation schedule. Select the "Now" radio button to generate a single report immediately. Select the "Periodically" radio button to create the report at specific times/intervals (an initial report will also immediately be generated). The "Simple" option enables you to schedule the report generation every "x" number of days, hours, minutes, seconds (e.g., every 5 days, every 5 minutes). The "Cron" option enables you schedule the report generation as a cron job (e.g., every minute, every hour, every year).
- **E-Mail** - Enter an address to e-mail an attached PDF of a report as configured in the Schedule field above. Each time a report is generated, an attached PDF of the report will be sent to the recipient. You can designate only one (1) e-mail recipient. Note that the E-Mail Preferences (Preferences - System Settings - E-Mail) must be configured for OmniVista to generate report e-mails.
- **Other Settings** - Click on this button to set optional report print parameters (e.g., page size, orientation). You can also add a description to the report.


2. After creating the report, go a supported application (e.g., Discovery, Locator, Analytics) and click on the **Add to Report** button at the top of the screen. The Add to Report Window will appear with the report/report view displayed in the Widget Name field (e.g., Inventory, Top N Ports Utilization Report Widget). Select the Report you configured in Step 1 from the Report Configuration drop-down list and click **OK**. A report for that screen will now be generated according to that report configuration.

You can generate reports for other applications based on the same report configuration by going to those applications and clicking on the Add to Report button and selecting the report from the Report Configuration drop-down list.

Note: The first time you configure a report (Step 1), a blank report is automatically generated and appears in the Report List. The report is blank because you have not yet associated the report with an application (Step 2). Once you complete Step 2, reports will be generated for that application based on the report configuration.


Note: You can also manually generate a report at any time by selecting the report and clicking on the **Generate Report** button on the Report Details Screen. You can only manually generate a report configured with the Schedule set to "Now". You cannot manually generate a report configured with a "Periodic" schedule.

Editing a Report

Select the report and click on the Edit icon  to bring up the Edit Report Configuration Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server. Note that you cannot edit the report title. You can edit the Report Settings, and/or click on the **Other Settings** button to


edit the print parameters. You can also remove a report from this Report Configuration by clicking on the "X" next to the field.

Deleting a Report

Select a report and click on the Delete icon , then click **OK** at the confirmation prompt. At the prompt, you have the option of deleting all reports associated with the report configuration. To delete them, select the "Also delete all generated reports" checkbox.

Report List

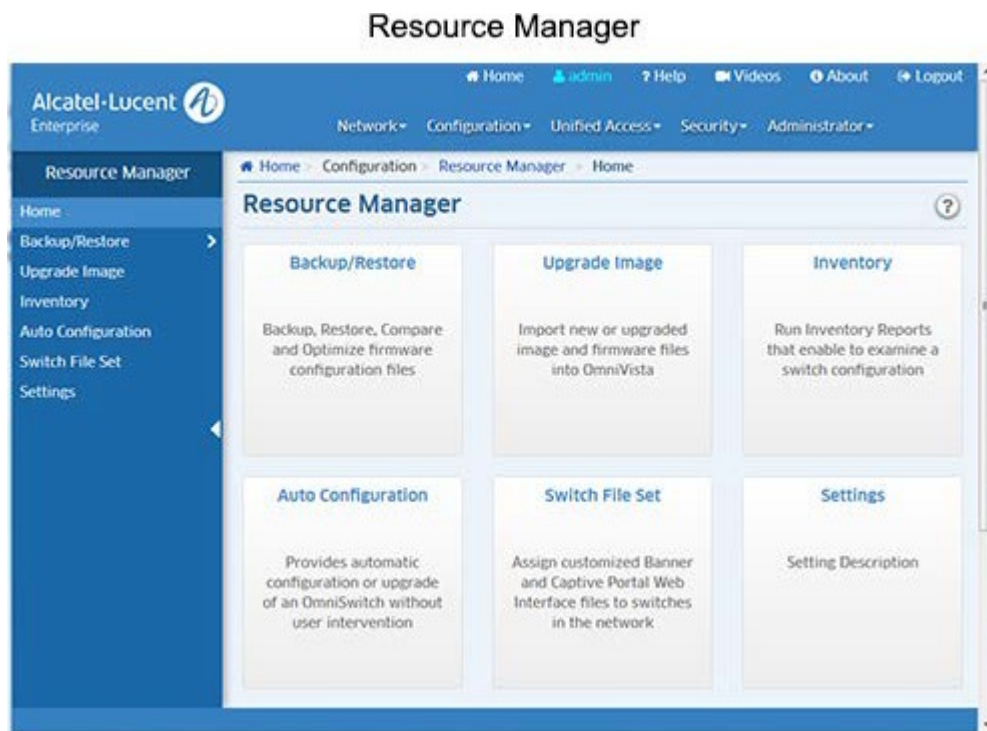
The Report List Screen displays all generated reports. Reports are displayed with the report creator's user name, report title, the date the report was created, and the version number of the generated report if applicable. For example, if a report was created by the "admin" user titled "Daily Users Report", on August 8, 2016, the report file name would be "admin_Daily Users_Report_20160808.pdf". Subsequent scheduled or manually-generated reports would have a version number added to the end of the filename (e.g., admin_Daily Users_Report_20160808_1.pdf, admin_Daily Users_Report_20160808_2.pdf).

To download/view a report in PDF format, select the report and click on the **Download** button. You can open the report for viewing or save the report. To delete a report(s), select the report(s) and click on the Delete icon , then click **OK** at the confirmation prompt. Note that you can only delete finished reports. You cannot delete a report in the "Generating" state.

21.0 Resource Manager

The Resource Manager application enables you to manage the firmware configuration files on network devices. Click on the applicable link, as described below, to carry out specific operations.

- **Backup/Restore** - Backup the current firmware configuration files in network devices to the OmniVista Server, and restore the configuration files to the devices when desired. You can also compare Configuration Backup Files on the same device or different devices, edit an existing backup and save the changes as a new backup file (AOS Devices), and view a summary of all stored backups.
- **Upgrade Image** - Import new or upgraded image and firmware files into OmniVista, and install the new files in network devices when desired. (Note that all new image files must be provided by Alcatel-Lucent Enterprise Customer Service.)
- **Inventory** - Run Inventory Reports on network devices that enable you to examine a device's configuration before performing the functions described above.
- **Auto Configuration** - Configure the Automatic Remote Configuration Feature. This feature provides automatic configuration or upgrade of an OmniSwitch without user intervention.
- **Switch File Set** - Assign customized Banner and Captive Portal Web Interface files to devices in the network.
- **Settings** - Used to set the amount of space that must be available on the CMM before an upgrade is allowed.



Backup

The Resource Manager Backup Screen displays a list of all backups that currently exist on the server. The screen used to backup configuration files for a network device. It is also used to schedule regular backups, edit a Configuration Backup File (boot.cfg), or delete a backup from the OmniVista Server. Backups can be used to restore configuration files to the network devices from which they were originally taken. You can also compare Configuration Backup Files on the same device or different devices, and view a summary of all stored backups.

Note: OmniVista supports the Multiple Working Directories Feature available on OS10K (AOS Release 7.2.1.R02 and later), OS6900 Switches (AOS Release 7.2.1.R01 and later), and OS6860 Switches (AOS Release 8.1.1.R01 and later). When performing Configuration-Only or Full Backup on

these devices, only the configuration or image files in the current running directory are backed up instead of the hard-coded “working” directory, in addition to the Certified Directory. The running directory can be any user-specified directory, including the Working Directory.

Performing a Backup

Click on the **Backup** button at the top of the screen to initiate a backup. Complete the screens as described below to backup one or more network devices. When you have completed all of the screens, click on the **Backup** button at the bottom of the screen to initiate the backup.

Backup Method

Select an option to choose a device selection method:

- **Backup by Devices** - To select specific devices from a list of discovered devices.
- **Backup by Maps** - To select a map(s) to backup all devices in the map(s). This option is used to backup **all** devices in the selected map(s). You cannot backup selected devices. To backup select devices, select the **Backup By Devices** option.

Note: If some devices in a map are not on-line, a dialog box will pop up warning you of the condition. Click **Yes** to continue the backup. Click **Cancel** to cancel the backup.

Devices/Maps Selection

The options on this screen will depend on the Backup Method selected above.

- **Backup by Devices** - Select the device(s) you want to backup from the list.
- **Backup by Maps** - Click on the Maps Selection button and select the map(s) containing the devices you want to back up.

Configuration

This screen is used to configure the type of backup performed (e.g., Full, Configuration Only, Images Only) and to schedule regular backups. Backup options (e.g., which directories to include, which files to include) vary according to the backup type. See the applicable section below for details on each backup type.

Full Backup

A Full Backup backs up both configuration files and image files. For AOS devices, all files in the Certified and Working directories are backed up. This includes all configuration-related files (user credentials, banner, time zone, etc.), and image files. If you are performing a Full Backup, select the directory(ies) to be backed up (Certified or All).

Note: AOS image files will not be FTPed from the switch. OmniVista will only record file version(s). Therefore, before Restore is to proceed, the required image file set must be stored in the Upgrade Image Repository. If the required images are not in the Repository, they will need to be imported using the Upgrade Image Screen in Resource Manager. Also note that if the image file info retrieved from the switch does not contain a file version, the file be physically copied from the switch.

Directory

If you are performing a Full Backup, select the directory(ies) to be backed up.

- **Certified** - Back up files in the Certified Directory.
- **All** - Back up files from the Working, Certified, Switch, and Network Directories.

Include Diagnostic and Dump Files

If you are performing a Full Backup on **all** files, you have the option to include/exclude Diagnostic and Dump Files. By default, Diagnostic and Dump files are not included in the backup. To include these files in the backup, set the **Include Diagnostic and Dump Files** slider to "On".

Description

Enter an optional description for the backup.

Configuration Only Backup

A Configuration Only Backup backs up all configuration-related files in all directories (including user credentials, banner, time zone, etc.). If are performing a Configuration Only backup, you will not have the option of selecting directories since all configuration-related files in all directories are backed up. However, on AOS devices you will have the option to include/exclude Security Files from the backup for security reasons.

Include Security Files

If you are performing a Configuration Only backup, you will not have the option of selecting directories since all configuration-related files in all directories are backed up. However, on AOS devices you will have the option to exclude Security Files from the backup for security reasons. By default, Security Files are included in the backup. If you do not want to include Security Files in the backup, set the **Include Security Files** slider to "Off".

Description

Enter an optional description for the backup.

Images Only Backup

An Images Only Backup backs up image files only. AOS image files will not be FTPed from the switch. OmniVista will only record file version(s). Therefore, before Restore is to proceed, the required image file set must be stored in the Upgrade Image Repository. If the required images are not in the Repository, they will need to be imported using the Upgrade Image Screen in Resource Manager. Note that if the image file info retrieved from the switch does not contain a file version, the file be physically copied from the switch.

Description

Enter an optional description for the backup.

Schedule Setting

Complete the applicable fields below to schedule a backup for a later time or to schedule regular backups.


- **Incremental** - If the **Incremental** option is set to "On", Resource Manager will initiate a scheduled backup only when there are changes in switch configuration files since the last Configuration only backup. The first incremental backup will be treated as a normal 'Configuration Only' backup and any successive incremental backups will be 'Configuration Only – Incremental' backups.
- **Start Time** - If you want to schedule a single backup at a later date and time, set the slider to "On" and set the date and time to initiate the backup.
- **End Time** - If you are scheduling regular backups (Simple or Cron), set the slider to "On" and set the date and time for the backups to end. The backups will occur as scheduled until the End Time.
- **Simple** - Select the Simple radio button to schedule regular backups at a specific interval (e.g., every day, every 7 days). Set the Repeat slider to "On" to repeat the schedule backup.
- **Cron** - Select the Cron radio button to create a regular backup cron job.

Review

The Review Screen enables you to review your backup configuration before initiating/scheduling the backup. If necessary, click on the **Back** button to make changes to the configuration. When you have verified the backup configuration, click on the **Backup** button to initiate/schedule the backup.

Note: If the CLI/FTP username and password for a device was not previously defined to OmniVista, you will be prompted to enter them before the backup can proceed.

Editing a Configuration Backup File

You can edit the contents of a Configuration Only Backup File (boot.cfg file). This file will be stored in the OmniVista Server and displayed in the Backup Table as "Backup" Select a Configuration Only Backup file in the Backup Table and click on the Edit icon . In the Select boot.cfg in drop-down menu, select the directory of the file you want to edit (working or certified) and click on the **Get** button. The contents of the file will appear in the File Content area. Edit the file and click on the **Save** button.

Deleting a Backup

Select the backup(s) you want to delete and click on the Delete icon . Click **OK** at the confirmation prompt.

Backup Information

The Backup Table displays basic information about all backups stored on the OmniVista Server. Click on a backup to view detailed backup file information.

Basic Information

- **Device Name** - The user-configured name of the device.
- **Device Address** - The IP address of the switch that was backed up.
- **Device Type** - The device/model type (e.g., OS6850E-P24X).
- **Date** - The date and time that the backup was initiated.
- **Backup Type** - The type of backup performed. The Backup type can be **Full Backup** (both configuration files and image files were backed up), **Configuration Only** (only configuration files were backed up), or **Image Only** (only image files were backed up).
- **Version** - The software version of the backup files (e.g., 6.4.6.186.R01).

- **Description** - The user-configured description for the backup, if applicable.

Detailed Information

- **Name** - The name of the individual file that was backed up and is currently stored on the OmniVista Server.
- **Directory** - The directory where the file was stored on the device (e.g., /flash/certified).
- **Version** - The firmware version of the file.
- **Description** - Alcatel-Lucent Enterprise provided description of the file.
- **Date** - The date the file was loaded into the switch.
- **File Size** - The size of the file, in bytes.

Important Facts About Backing Up

When performing a backup, firmware configuration files are FTPed from the switch to the OmniVista server. To gain access to the switch, the FTP user name and password must be known to OmniVista. You can specify FTP user names and passwords via the Edit Discovery Manager Entry window. (See the Topology help for further information.) If you did not define FTP Logic names and passwords via the Edit Discovery Manager Entry window, and you attempt to save or restore configuration files, you will be queried for the FTP username and password for each individual device for which files are being saved or restored. If the FTP username and password are not supplied to OmniVista, the FTP process will return errors and the device will not be backed up. The process of backing up other switches will continue. Firmware is automatically copied and restored via FTP, and any errors that can occur when performing these tasks outside of OmniVista are also possible when using OmniVista.

If a backup operation fails in the middle of the backup operation (which could occur if a device goes down between the server and the target switch), no files are saved on the server. If the full complement of files are not saved, any initial files that were saved are deleted from the server.

Backups of AOS devices include the contents of the certified directory and working directory. Only files in the flash memory of the primary MPM module are saved. No files are saved that end with .err , .dmp , /.. , or /. , as these files are either temporary or will cause problems during the FTP process due to conflict with system file names.

Important Notes: The configuration files saved are those in flash memory and are not necessarily the configuration files that the switch is currently running. The files are not zipped to save disk space on the OmniVista Server. The user may perform multiple backups on the same day, if so desired.

Users should not attempt to copy configuration files saved on the OmniVista Server to other machines. The saved files contain binary configuration information, including the IP address/MAC address of the source machine, and using these files on another machine could bring the network down.

Note: SFTP will be used when a device is configured in OmniVista to use SSH. If a device is configured to use SSH in OmniVista, SSH must be enabled on the device itself.

Restore

The Resource Manager Restore Screen displays a list of all device backups and is used to restore the configuration to the device from a previous backup. You can only restore the configuration to the original device from which the backup was taken. (Backups cannot be restored to other switches, because doing so would cause mismatched IP addresses and other network problems.)

Performing a Restore

Files Selection

Click on the **Add/Remove Backup Files** button. (If you are performing the restore from the Master Dashboard, select a customer network from the drop-down list.). A list of devices that were backed up is displayed in the Backup File Selection window. Select the device(s) on which you want to perform the restore and click **OK**. Note that you can click on a device to view a list of all of the files that will be restored. By default, all of the backed up files are selected and will be restored. However, you can de-select files you do not want to include in the restore. Click on the **Next** button to go to the Configuration Screen.

Configuration

AOS Switch(es) Only

If the selected switch(es) is an **AOS 6.x Switch**, and you selected files only in the Certified Directory, you have the option of restoring files to the Working Directory or storing the files to both the Working and Certified Directories. Select the **Restore to Working Directory** or **Restore to Working & Certified Directory** radio button to specify the directory(ies) to which you want the backup restored. If you select files in only the Working Directory or in both the Certified and Working Directories, you will not have this option. The files will be restored to their respective directories. Also note that this option is not available for AOS 7.x or 8.x switches.

Note: OmniVista supports the Multiple Working Directories Feature available on OS10K (AOS Release 7.2.1.R02 and later), OS6900 Switches (AOS Release 7.2.1.R01 and later), and OS6860 Switches (AOS Release 8.1.1.R01 and later). When performing Configuration-Only or Full Backup on these devices, only the configuration or image files in the current running directory are backed up instead of the hard-coded “working” directory, in addition to the Certified Directory. The running directory can be any user-specified directory, including the Working Directory.

Option

Select the options to be taken if the following changes are detected on the device:

- **Continue to restore when chassis has changed** - Select this option if you want to continue the restore even if it is found that the chassis contents, or the chassis type, has changed since the backup. If you do not enable this checkbox, the restore will not take place if the chassis has changed.
- **Continue to restore when detected new image files** - Select this option if you want to continue the restore even if it is found that a new image file resides on the device (i.e., a file that was not previously backed up). If you do not enable this checkbox, the restore will not take place if a new image file is found on the device.

Click on the **Restore** button then click **Yes** at the Confirmation Prompt to initiate the restore. When the restore has successfully completed, click on the **Go to Topology to Reboot Device** link. The Topology application will open with the device(s) highlighted. Click on **Reboot** in the Device Operations area to reboot the device(s) to load the restored configuration into flash memory.

Note: You **must** reboot the device(s) to complete the restore operation.

Compare

The Resource Manager Compare Screen enables you to compare Configuration Backup Files on the same device or different devices using a "Diff" Utility to view any differences between the files on a line for line basis. You can compare files on different devices or compare files on the same device . You can also use the utility to compare text files on the local file system .

Note: The "boot.cfg" file is the target of this utility, however you can use it to compare any text- based files. You cannot use the utility to compare any binary files (e.g., .img, jpg, jar).

Selecting Files

The File Diff Screen is used to select the files you want to compare. To compare backup files from different devices or backups from the same device, select "Backup File" from the **Select From File** drop-down menu on the left side of the screen. Click on the **Browse** button to bring up a list of current backups. Select a backup to bring up a list of files contained in the backup. Select the file you want to compare, and click **OK**. Repeat the steps to select a backup file on the right side of the screen. When you have selected both files, click on the **Compare** button. The file comparison is displayed in the File Diff Window.

Note: To compare text files on the local file system, select "Local" from the from the **Select From File** drop-down menu. Click on the Browse button and browse to the files on the local system.

Comparing Files

The File Diff Window displays the files side-by-side with all of the differences highlighted (Changed, Inserted, Deleted). You can use the Arrow keys at the top of the screen on the right side of the window to jump to each change; or you can select a specific change from the **Select to Jump** drop-down menu. You can also use the scroll bars to scroll through the documents and view changes.

Backup Summary

The Resource Manager Backup Summary Screen provides a status summary (Successful, Failed) of all backups stored on the OmniVista Server.

- **Device Address** - The IP address of the switch that was backed up.
- **Last Backup Attempt Date** - The date and time that the backup was initiated.
- **Last Successful Backup Date** - The date and time that the backup completed, if successful. If the backup failed, the field will be blank and all other fields will be highlighted in red.
- **Date** - The date and time that the backup was initiated.
- **Backup Type** - The type of backup performed. The Backup type can be **Full Backup** (both configuration files and image files were backed up), **Configuration Only** (only configuration files were backed up), or **Image Only** (only image files were backed up).
- **Status** - The backup status (BACKUP_SUCCESS, BACKUP_FAIL). If a backup fails, all fields will be highlighted in red.
- **Description** - The user-configured description for the backup, if applicable.
- **Message** - The backup status message. Indicates whether or not the backup was successful (Backup Successful) or failed. If "failed" additional information is provided.

Upgrade Image

The Resource Manager Upgrade Image Screen displays all of the Software and Firmware Files stored in the Upgrade Image Repository on the OmniVista Server. These files are used to upgrade software, firmware, and FPGA files on network devices. Once you download the files from the Customer Support Web Site, you can import the files into the Upgrade Image Repository, and install the upgrade software and firmware on devices on the network. Note that FPGA upgrade is only supported on OS9000, OS6450, and OS6250 Switches running AOS 6.6.4.R01 and later.

Note: OmniVista supports the Multiple Working Directories Feature available on OS10K (AOS Release 7.2.1.R02 and later), OS6900 Switches (AOS Release 7.2.1.R01 and later), and OS6860 Switches (AOS Release 8.1.1.R01 and later). On these devices, the Upgrade operation installs the files to the user-specified directory instead of the hard-coded Working Directory.

CAUTION: Never attempt to import or install firmware files or upgrade packages acquired from any source other than Alcatel-Lucent Enterprise Customer Service. Image and Firmware files are specially packaged by Alcatel-Lucent Enterprise Customer Service for import into OmniVista, and contain an LSM file that describes the package contents to OmniVista. Resource Manager will prevent unsupported upgrades. When such an attempt is made, an error message is displayed informing the user that the upgrade has been rejected. This message also displays details of the versions of the switch software required to successfully perform the upgrade.

WARNING: If you are performing an image file upgrade **and** a U-Boot/Miniboot upgrade, you **must complete the image file upgrade before** upgrading the U-Boot and Miniboot files.

Importing the Upgrade Files

All upgrade files supplied by Alcatel-Lucent Enterprise Customer Service are packaged as WinZip executables and have a *.zip file extension. Do not attempt to unzip the firmware files manually. When you Import the WinZip executable, OmniVista will automatically unzip the executable as part of the import process. Once the file is imported, the File Set (which contains all of the individual files) appears in the File Sets Table.

Upgrade files are available on the Alcatel-Lucent Enterprise Customer Service website. Download the file to your PC from the website. After downloading the file, click on the **Import** button to locate and import the file to OmniVista.

Note: You can delete a file from the File Sets Table by selecting the file(s), clicking on the delete icon , then clicking **OK** at the Confirmation Prompt.

Installing the Upgrade Files

Remember, if you are performing an image file upgrade **and** a U-Boot/Miniboot upgrade, you **must complete the image file upgrade before** upgrading the U-Boot and Miniboot files. Select a File Set in the File Sets Table and click on the **Install** button. The Install Upgrade Image Software Wizard guides you through the upgrade process. Each screen in the wizard is detailed below.

Note: The switch FTP timeout default is 5 minutes, so the upgrade will fail if the time to transfer files from OmniVista to a switch is over 5 minutes. It is recommended that you increase the FTP timeout in switches you are upgrading to a higher value to make sure there is enough time to transfer files (CLI command: **session ftp timeout <time>**).

Firmware File Selection

All of the files in the File Set that you selected are displayed in the File Detail area. (The name of the imported Zip File is displayed in the File Name field.) Select the file(s) you want to install and click **Next** to go to the Devices Selection Screen.

Devices Selection

All devices that qualify for installation of the selected files are displayed in the Select Devices area. Select the device(s) in which you want to install the file(s), then click **Next** to go to the Software Installation Screen.

Note: The lowest supported AOS version for OS6855P-14 devices is AOS 6.4.4.9.R01. The lowest supported AOS version for OS6450 devices is AOS 6.6.3.360.R01. The highest supported AOS version for OS6400 devices is AOS 6.4.5.487.R02.

Software Installation

Select the installation options as described below.

Installation Options

- **Upgrade BMF Files** - Upgrade the BootROM, MiniBoot, or FPGA files (AOS switches only).
- **Upgrade Images Files** - Upgrade the image files on the switch(es) (Default) .
- **U-Boot Upgrade on all NIs** - Perform u-boot upgrade for all the NIs on the switch(es) (9000 series switches only).
- **In-Service Software Upgrade (ISSU)** - Upgrade the image files on redundant CMMs with minimal data interruption. This option is available (and displayed) for OS10K and OS6900 (7.3.1.R01 and later) and OS6860 (8.1.1.R01 and later). ISSU support for the OS10K is for both standalone and virtual chassis; ISSU support for the OS6900 and OS6860 is for Virtual Chassis configuration only. [Click here for more information on ISSU.](#)
- **Directory** - The Directory field is enabled when the Upgrade Image Files checkbox is selected and the images are for devices supporting the Multiple Working Directories Feature - OS10K (7.1.1.R01 and later), OS6900 (7.2.1.R01 and later) OS6860 - 8.1.1.R01 and later). The directory path must be either an absolute path (e.g. /flash/myimagedir) or a relative path to the flash (“/flash/” will be prefixed in this case). Validations will be done to ensure the directory path is valid before the images are sent to the switches. Note that if the user-specified directory does not exist on the switch, it is automatically created. Once the images are uploaded to the switch, if the user-specified directory does not contain any boot.cfg file, it is copied from the current running directory of the switch.

Note: The ISSU upgrade procedure for upgrading AOS from 8.1.1.xxx to 8.2.1.304.R01 on OS6860 and OS6860E Switches is different than the regular ISSU upgrade procedure. OmniVista does not support this ISSU upgrade path, please refer to APPENDIX C of the 8.2.1.304.R01 Release Notes for detailed instructions on the upgrade procedure.

6200 Devices Options

- **6200 Series Installation Options** (6200 Devices Only). If you are upgrading a stack of devices, the following options will be enabled.
 - **Upgrade Master Unit Only** - Upgrade the image files on the master switch in the stack.
 - **Upgrade All NIs in Stack** - Upgrade the image files on all switches in a stack.

When you have configured all of the applicable installation options, click on the **Install Software** button to initiate the upgrade process. When the install has successfully completed, click on the **Go to Topology to Reboot Device** link at the top of the screen. The Topology application will open with the device(s) highlighted. Click on **Reboot** in the Device Operations area to reboot the device(s) to load the restored configuration into flash memory.

Note: You **must** reboot the device(s) to complete the upgrade. Remember that image files are installed into the working directory of AOS devices. After the installation completes, you should reboot AOS devices. You may also want to save the working directory to the certified directory.

ISSU Upgrade

The In-Service Software Upgrade (ISSU) feature is used to upgrade the CMM images running on supported devices with minimal disruption to data traffic. The CMM images can be upgraded only on fully synchronized, certified, and redundant systems. A minimum of size of mandatory images + 3MB flash space must be present in the device to accommodate the image files that are used to upgrade existing image files. The ISSU upgrade process is the same as the upgrade process detailed above. However, you cannot select individual files from the File Set. All of the files will be installed. You cannot select individual files in the File Details area. The following CMM images are ISSU capable.

Prior to FTPing the images to switches, Resource Manager performs the following checks to make sure the selected device is ready for ISSU:

- Ensures the device is redundant, fully certified, and synchronized.
- Ensures that sufficient flash space is available on the primary CMM (a minimum of size of mandatory images + 3MB flash file system space is required for the upgrade).

Note: Although Resource Manager will make certain that the switch is ISSU capable, it will not perform any check whether selected ISSU images are compatible with the particular software version running on the switch. This information will be provided to customers by Customer Support when a new ISSU package is released.

If any of these checks fail for a device, Resource Manager logs the error message, and continues with the next device. Otherwise, Resource Manager checks for the existence of the /flash/issu directory on the primary CMM, and creates the directory if it is not present. If the directory already exists and is not empty, Resource Manager removes all files in the directory before replacing them with the new images.

When Resource Manager finishes issuing the ISSU command to the selected devices, the user is asked to perform "Copy Working to Certified" for each device. Make sure you also perform: "Flash-Synchro" for each device.

ISSU Upgrade Paths

AOS Release 7¶

	Upgrading From 7.3.4.R02 GA or 7.3.4.R02 Maintenance Release	Upgrading from any other 7.X Release
OS6900 - VC	ISSU - Supported Standard Upgrade - Supported	ISSU - Not Supported Standard Upgrade - Supported
OS6900 - Standalone	ISSU - N/A Standard Upgrade - Supported	ISSU - N/A Standard Upgrade - Supported
OS10K - VC	ISSU - Supported Standard Upgrade - Supported	ISSU - Not Supported Standard Upgrade - Supported
OS10K - Standalone (Dual-CMM)	ISSU - Supported Standard Upgrade - Supported	ISSU - Not Supported Standard Upgrade - Supported
OS10K - Standalone (Single-CMM)	ISSU - N/A Standard Upgrade - Supported	ISSU - N/A Standard Upgrade - Supported

AOS Release 8

	Upgrading From 8.2.1.R01 GA or 8.2.1.R01 Maintenance Release	Upgrading from any other 8.X
OS6860-VC	ISSU - Supported Standard Upgrade - Supported	ISSU - Not Supported Standard Upgrade - Supported
OS6860-Standalone	ISSU - N/A Standard Upgrade - Supported	ISSU - N/A Standard Upgrade - Supported

Important Information About Upgrades

If FTP User Names/Passwords are Undefined

If the FTP user names and passwords for the devices were not previously defined to OmniVista, the FTP User Name Password window displays. To supply the FTP user name and password for a device, select the device in the FTP User Name Password window and click the **Edit** button. Enter the FTP user name and password for the selected device in the appropriate fields. If the user name and password you enter also apply to the other devices, click the **Same for all Unspecified** checkbox. Then click the **OK** button.

If necessary, continue to enter FTP user names and passwords until they have been specified for all devices listed. When all user names and passwords have been specified, click Yes at the installation confirmation prompt to initiate the installation process.

If Version Numbers are Older

If the image files being installed have an older version number (or the same version number) than the image files currently resident on a device, a warning message will appear. Note that installing older versions of image files may result in a loss of functionality or, in the case of OmniStack devices, the resetting of device parameters. Click **Yes** if you wish to perform the installation anyway. Click **No** to cancel the installation.

File Sets Information

The individual upgrade files are contained in File Sets for each device type. The File Sets Table displays all of the File Sets stored in the Upgrade Image Repository on the OmniVista Server. The table displays basic information. Click on a file set to display detailed information for the files contained in the File Set.

Basic Information

- **Type** - The upgrade file set device type (e.g., Omniswitch6860).
- **Date** - The date the file set was imported into OmniVista.
- **Version** - The firmware version of the file set (e.g., 8.2.1.304.R01).
- **Description** - The file set description (e.g., AOS 8.2.1.304.R01).

Detailed Information

- **File Name** - The name of the file (e.g., Uos.img).
- **Version** - The firmware version of the file set (e.g., 8.2.1.304.R01).
- **Description** - The file description (e.g., Alcatel-Lucent OS).
- **Date** - The date the file was created.
- **File Size** - The file size, in bytes.

Inventory

The Resource Manager Inventory Screen is used to create Inventory Reports on network devices that enable you to examine a device's configuration. A Switch Inventory Report includes system information, detailed module information, chassis data, and health information for an individual switch. You can request an Inventory Report for a single device or for multiple devices simultaneously.

Creating an Inventory Report

Select the type of report you want to create from the Report Type drop-down menu:

- **Condensed Content** - Generates a smaller, condensed report that is displayed when it is generated. If you select this option, click on the **Browse** button and select the device(s) for which you want to create the report, then click on the **Create** button.
- **Detailed Content** - Generates a larger, detailed report. This report is not immediately displayed. The report is generated as an HTML file and a link is displayed to access the report. The report is available for two (2) days before it is purged. If you select this option, enable any specific information that you want to include in the report (e.g., System Information, Chassis Information, 8021AB Information). Click on the **Browse** button and select the device(s) for which you want to create the report, then click on the **Create** button.

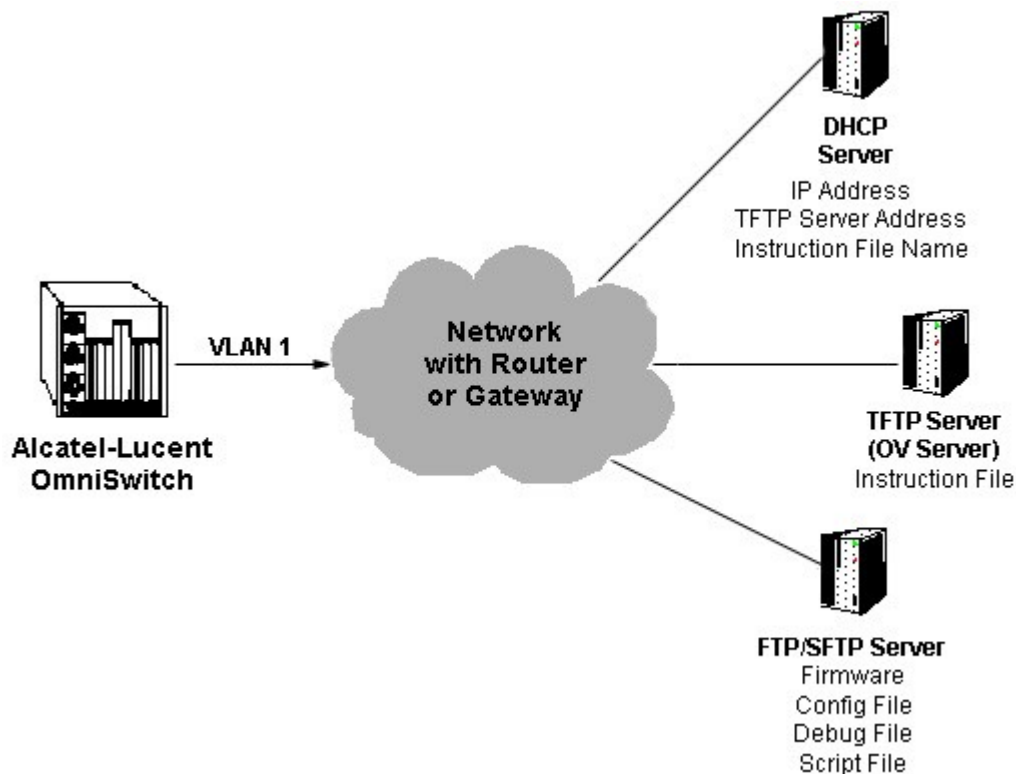
Auto Configuration

The Resource Manager Auto Configuration Screen is used to configure the Automatic Remote Configuration Feature. This feature provides automatic configuration or upgrade of an OmniSwitch without user intervention by pushing an Instruction File to the device. The Instruction File contains all of the information required to automatically locate and download all of the necessary files to configure a new device/upgrade an existing device on the network. When a device is initially deployed in a network, the Instruction File is sent to the device to download the applicable Image, Configuration, Debug, and Script Files from remote servers to bring the device online in the network. The Auto Configuration Screen displays all configured Instruction Files. It is also used to create, edit, and delete Instruction Files.

Auto Configuration Overview

The Auto Configuration Feature automatically configures a new switch and brings it online in the network. In addition, the feature can be used to automatically upgrade a switch with new Firmware, Configuration, and Debug files. As shown below, the Auto Configuration Feature requires a Default DHCP Server, a TFTP Server (the OmniVista Server) that contains the Instruction File, and a remote FTP/SFTP Server that contains the Firmware, Configuration, Debug, and Script Files.

- **DHCP Server** - Provides the switch with an IP address as well as the location of the TFTP Server and the name of the Instruction File. The switch must have at least one port with connectivity to the DHCP Server through Default VLAN 1.
- **TFTP Server** - Resides on the OmniVista Server and contains the Instruction File, which contains the file names and locations of the Firmware, Configuration, Debug, and Script Files stored on the FTP/SFTP Server. The OmniVista 2500 TFTP Server Code Library transfer limit file size is 4GB, per RFC 2347.
- **FTP/SFTP Server** - Contains the Firmware, Configuration, Debug, and Script files. (A Remote FTP/SFTP Server is optional. You can also store the files on the TFTP Server.)



Auto Configuration on a New Switch

New OmniSwitches are shipped without a boot.cfg file. When the new switch is connected to the network as a new device with no boot.cfg file in the working directory, the Automatic Remote Configuration process is initiated. First, a DHCP client is automatically created on VLAN 1 on the switch and the switch obtains an IP address as well as the address of the TFTP Server and the name of the Instruction File. The switch then downloads the Instruction File. The Instruction File contains the file names and locations of the Firmware, Configuration, Debug, and Script Files which are then downloaded from the FTP/SFTP Server and saved as the boot.cfg file in the /flash/working directory. The DHCP Client on VLAN 1 is removed and the Script File is launched to configure the switch and the switch is automatically rebooted to load the image files from the /flash/working directory.

Note: You must create an Instruction File for each switch model on your network. When the switch sends its initial request to the DHCP Server, the model name (e.g., 6850, 9000) is included in the Vendor Class Identifier Field (Option 60 Field). The DHCP Server will then return the Instruction File corresponding to the model listed in the field.

Automatic Configuration Updates

In addition to automatically configuring new switches on the network, the Auto Configuration Feature can be used to automatically update existing network switches upon reboot. To enable a switch to be automatically configured with the latest image files and configuration files on reboot, remove the boot.cfg file from the /flash/working directory. When the switch reboots and the network detects that there is no boot.cfg file, the Automatic Remote Configuration process is initiated. The Automatic Configuration software compares the current firmware version on the switch with the version stored on the FTP/SFTP Server. If the version on the switch is older than the version on the FTP/SFTP Server, the Automatic Configuration process is launched, as described above.

Automatic Configuration Files

The files downloaded during the Automatic Configuration Process are detailed below.

- **Instruction File** - The initial file required for the automatic remote configuration process to occur. The file contains the names and location of the Firmware, Configuration, Debug and Script files, which are stored on a remote FTP/SFTP Server.
- **Firmware Files** - Image files that are used to initially configure or upgrade a switch. The firmware files, which differ for different OmniSwitch platforms, contain the executable code, which provides support for the system, Ethernet ports, and network functions.
- **Configuration File** - Bootup configuration information for the switch (network configuration parameters).
- **Debug Configuration File** - Default debug configuration.
- **Script File** - This file contains the commands to be performed on the switch so that appropriate actions can be taken on the downloaded files (Firmware, configuration and Debug Files). The Script File can be created using CLI commands, which are performed in the order in which they appear in the script. A Script File example is shown below:

```
reload working no rollback-timeout
copy working certified flash-
synchro
```

Note: If a 'write memory' command is used in the script file, it overwrites the boot.cfg file. The Script File should not contain the write memory command if it is downloaded along with the configuration file.

For more information on configuring the Script File, See the "Managing Automatic Remote Configuration" chapter in the *Network Configuration Guide* .

Quick Steps for Automatic Remote Configuration

The steps below give a quick overview of configuring a switch for Auto Configuration. Follow the steps below to configure Automatic Remote Configuration for you network. Detailed instructions, including Script File Syntax and examples can be found in the "Managing Automatic Remote Configuration Download" chapter of the *Switch Management Guide*.

Note: The switch must have at least one port with connectivity to the DHCP Server through default VLAN 1.

1. Configure the default network DHCP Server with the TFTP Server address (Option 66) and Instruction File name (Option 67). For example

- **Option 66:** 128.251.17.224 (TFTP Server address. This is the OmniVista Server address)
- **Option 67:** os6900.alu (name of the Instruction File on the OmniVista Server).

Note: For details on how to configure the DHCP server, see the "Configuring DHCP Server" chapter in the *Network Configuration Guide* .

2. Configure a Script file. See the "Managing Automatic Remote Configuration Download Chapter" in the *Network Configuration Guide* .

3. Store the Firmware, Configuration, Debug and Script Files on the FTP/SFTP Server.

4. Create the Instruction File.

Creating an Instruction File

The Instruction File contains all of the information needed by a device to locate and download the applicable Image, Configuration, Debug, and Script Files from remote servers. The Firmware, Configuration, Debug and Script files, differ for different OmniSwitch platforms. You must create an Instruction file for each switch model on your network (e.g., 9000, 6850, 6855) . When a new switch comes online, the switch type is sent to the DHCP Server using Option 60 to select the Instruction File for that device type. To create an Instruction File, click on the Create icon and complete the Instruction Fields as described below.

Header/File Servers


- **Instruction File Path** - The Instruction File directory path. The Instruction File is stored on the OmniVista Server in the root\data\resourceanager\instructionfiles\ directory. Enter the directory of the Instruction File on the Server (e.g., os6855).
- **Instruction File Name** - The Instruction File name. You can create multiple Instruction Files (e.g., Instruction Files for different models - 9000, 6850, 6855). When a new switch comes online, the switch model is sent to the DHCP Server using in the Vendor Class Identifier Field. The DHCP Server will then return the TFTP Server Address and Instruction File path for the corresponding Instruction File based on the information configured in the Option 66 Field. You must use the ".alu" extension for any Instruction Files you create (e.g., instruction1.alu).
- **Instruction File Header** - User-configured header for the Instruction File. This may contain any user information such as switch ID, file version, etc.
- **Primary File Server**

- **Primary Server IP Address** - The IP address of the Primary FTP/SFTP Server.
- **Primary Server Protocol** - The protocol used to communicate with the Primary Server (FTP, SFTP).
- **Primary Server User** - The user name of the primary user (e.g., admin).
- **Secondary File Server** (Optional)
 - **Secondary Server IP Address** - The IP address of the Secondary FTP/SFTP Server. If OmniVista is unable to connect to the Primary Server after three (3) retries, OmniVista logs the error and connects to the Secondary File Server. A Secondary Server is not required. If you do not want to add a Secondary Server, make sure all of the Secondary Server fields are empty. The IP Address and User Fields must be empty and the Protocol Field must be set to "None".
 - **Secondary Server Protocol** - The protocol used to communicate with the Secondary Server, if applicable (e.g., FTP). If a Secondary Server is not configured, this field must be set to "None".
 - **Secondary Server User** - The user name of the secondary user (e.g., admin). If a Secondary Server is not configured, this field must be empty.


Software and Config Files

The only required fields in this section are the Firmware Version and Firmware Location fields. The remaining fields are not required and the fields can be left empty.


- **Firmware Version** - The version of the firmware to be downloaded from the FTP/SFTP Server (e.g., OS_6_4_6_101_R01). You must use the format shown in the example.
- **Firmware Location** - The directory location of the firmware on the FTP/SFTP Server (e.g., /ftproot/firmware).
- **Config File Name** - The name of the Configuration File (e.g., boot.cfg).
- **Config Location** - The location of the Configuration File on the FTP/SFTP Server (e.g., /ftproot/config).
- **Debug File Name** - The name of the Debug File on the FTP/SFTP Server (e.g., AlcatelDebug.cfg).
- **Debug Location** - The location of the Debug File on the on the FTP/SFTP Server (e.g., /ftproot/debug).
- **Script File Name** - The name of the Script File on the FTP/SFTP Server (e.g., OS6850_script.txt). If a script file is not specified in the Instruction File, or if it is not properly downloaded, the Automatic Remote Configuration Manager software automatically initiates a "reload working no rollback-timeout" command after firmware or bootup configuration files are downloaded.
- **Script File Location** - The location of the Script File on the FTP/SFTP Server, if applicable (e.g., /ftproot/script).
- **License File Name** - The name of the License File on the FTP/SFTP Server.
- **License File Location** - The location of the Script File on the FTP/SFTP Server, if applicable (e.g., /ftproot/ license).

Note: Since many of the fields for different Instruction Files will be the same (e.g., File Server Address, Firmware Location), a shortcut to creating additional Instruction Files is to select an existing file in the Instruction File List, click on the Create icon , and change only the fields that are different for the new file (e.g., Instruction File Path, Instruction File Name).

Editing an Instruction File

Select the file in the Instruction Files List and click on the Edit icon . Edit the necessary fields as described above, click on the **Apply** button, then click **OK** at the confirmation prompt to update the Instruction File on the TFTP Server.

Deleting an Instruction File

Select the file(s) in the Instruction Files List, click on the Delete icon , then click **OK** at the confirmation prompt.

The Instruction Files List

The Instruction Files List displays information about all Instruction Files stored on the OmniVista Server.

- **Instruction File Path** - The Instruction File directory path. The Instruction File is stored on the OmniVista Server in the root\data\resource\manager\instructionfiles\ directory. Enter the directory of the Instruction File on the Server (e.g., os6855).
- **Instruction File Name** - The Instruction File name. You can create multiple Instruction Files (e.g., Instruction Files for different models - 9000, 6850, 6855). When a new switch comes online, the switch model is sent to the DHCP Server using in the Vendor Class Identifier Field. The DHCP Server will then return the TFTP Server Address and Instruction File path for the corresponding Instruction File based on the information configured in the Option 66 Field. You must use the ".alu" extension for any Instruction Files you create (e.g., instruction1.alu).
- **Instruction File Header** - User-configured header for the Instruction File. This may contain any user information such as switch ID, file version, etc.
- **Primary Server IP Address** - The IP address of the Primary FTP/SFTP Server.
- **Primary Server Protocol** - The protocol used to communicate with the Primary Server (FTP, SFTP).
- **Primary Server User** - The user name of the primary user (e.g., admin).
- **Secondary Server IP Address** - The IP address of the Secondary FTP/SFTP Server. If OmniVista is unable to connect to the Primary Server after three (3) retries, OmniVista logs the error and connects to the Secondary File Server. A Secondary Server is not required. If you do not want to add a Secondary Server, make sure all of the Secondary Server fields are empty. The IP Address and User Fields must be empty and the Protocol Field must be set to "None".
- **Secondary Server Protocol** - The protocol used to communicate with the Secondary Server, if applicable (e.g., FTP). If a Secondary Server is not configured, this field must be set to "None".
- **Secondary Server User** - The user name of the secondary user (e.g., admin). If a Secondary Server is not configured, this field must be empty.
- **Firmware Version** - The version of the firmware to be downloaded from the FTP/SFTP Server (e.g., OS_6_4_6_101_R01). You must use the format shown in the example.
- **Firmware Location** - The directory location of the firmware on the FTP/SFTP Server (e.g., /ftproot/firmware).
- **Config File Name** - The name of the Configuration File (e.g., boot.cfg).
- **Config Location** - The location of the Configuration File on the FTP/SFTP Server (e.g., /ftproot/config).
- **Debug File Name** - The name of the Debug File on the FTP/SFTP Server (e.g., AlcatelDebug.cfg).
- **Debug Location** - The location of the Debug File on the on the FTP/SFTP Server (e.g., /ftproot/debug).
- **Script File Name** - The name of the Script File on the FTP/SFTP Server (e.g., OS6850_script.txt). If a script file is not specified in the Instruction File, or if it is not properly downloaded, the Automatic Remote Configuration Manager software automatically initiates a "reload working no rollback-timeout" command after firmware or bootup configuration files are downloaded.
- **Script File Location** - The location of the Script File on the FTP/SFTP Server, if applicable (e.g., /ftproot/script).
- **License File Name** - The name of the License File on the FTP/SFTP Server.
- **License File Location** - The location of the Script File on the FTP/SFTP Server, if applicable (e.g., /ftproot/license).

Switch File Set

The Resource Manager Switch File Set Screen is used to create a command prompt Login Banner and/or Captive Portal Web Page file and assign the file to devices on the network. A Banner file is a .txt file that is displayed when a user first logs into a network device using the command line interface. Banner files can be customized to display a unique command line banner for all devices on the network. Captive Portal, a web-based user authentication option within the Access Guardian application. A Captive Portal file is an HTML file that is presented to the user with a web page for authentication. A Switch File Set contains Banner or Captive Portal files that can be assigned to network devices to be presented to users when they login to a device.

The Switch File Set Screen displays all configured Switch File Sets and is used to create, edit, delete, and assign Switch File Sets.

Note: Before assigning Banner or Captive Portal files to all devices in the network, it is recommended that you customize the file(s) and send the file(s) to a single switch on the network for verification. When you are satisfied with the customized file(s), you can then push the files to the network. Any subsequent changes to the files can be made on that same switch, and the new files imported and pushed to the network.

Overview

Below are a list of Banner and Captive Portal file names. These files are stored in OmniVista and will be the Banner and Captive Portal Web files (e.g., Login Page, Help Pages) that you will customize for your network. The files you create must use these file names. For example, if you create a Captive Portal Login Page, the file must be named *cpLoginWelcome.inc*. Once you have created all of the necessary files and verified them on a network device, you can then import those files from that device and "push" them other devices on the network. The file names and their use are described below.

- **banner.txt** - A Banner file is a .txt files that is displayed when a user first logs into a network device using the command line interface.
- **background.gif/.jpg/.png** - Use this file to provide a page background image that Captive Portal will display on all pages.
- **cpLoginHelp.html** - Use this file to customize the Captive Portal login help page. A question-mark ("?") button links to this HTML help page, which is displayed in a separate browser window
- **cpLoginWelcome/cpStatusWelcome/ cpFailWelcome/cpBypassWelcome.inc** - Use these files to customize the welcome message for the Captive Portal login, successful status, fail status, and bypass status page.
- **cpPolicy.html** - The User Acceptable Policy HTML file that is linked to the Captive Portal login page. The link provided opens a new browser window to display the policy information.
- **logo.gif/.jpg/.png** - Use these files to provide a company logo that Captive Portal will display on all pages.

Note: Create custom logo and background pages using the .gif, .jpg, or .png formats. Captive Portal checks the flash/switch directory on the switch for a .gif file, then a .jpg file, and finally a .png file. Whichever file type Captive Portal encounters first is the file used to display the custom logo or background.

The .inc files, which are used to present customized welcome messages, are partial HTML files that can include only text or text and other HTML tags, such as links. Note that these .inc files are wrapped in a paragraph HTML tag within the body of a Captive Portal default page.

Banner Files

A Banner file is a .txt file that is displayed when a user first logs into a network device using the command line interface (e.g., a company name, device name). You must first create the file, then assign ("push") the file to devices on the network.

Captive Portal Files

Captive Portal is a configurable option within the Access Guardian application that allows web-based clients to authenticate through 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.

Creating a Switch File Set

A Switch File Set contains Banner or Captive Portal files that can be assigned to network devices to be presented to users when they login to a device. To create a Switch File Set, click on the Create icon to bring up the Create Switch File Set window. Enter a **File Set Name** and **File Set Description** and select the **File Set Type** from the drop-down menu (Captive Portal/Banner). All of the default files for that File Set Type are displayed in the Files Table. Select the file(s) you want to include in the Switch File Set and click on the **Create** button.

You can add custom files to the Switch File Set by adding files from your PC or importing the files from a network device. To add files from your PC, click on the **Add** button, locate the file(s), and click **OK**. To import files from a network device, click on the Import icon, select the switch from which you want to import the files, and click on the **Import** button. The files you add or import will appear in the Files Table. You can then select that file (or files) and add them to the Switch File Set. Remember, the files you add must have the same file name (e.g., banner.txt).


Once you have created the Switch File Set, you must assign it to device(s) on the network.

Assigning a Switch File Set


Select the Switch File Set that you want to assign from the Switch File Set Table and click on the **Assign** button. Select the device(s) to which you want to assign the Switch File Set, click the **Apply** button, then click **OK**.

Note: Before assigning Banner or Captive Portal files to all devices in the network, it is recommended that you customize the file(s) and send the file(s) to a single switch on the network for verification. When you are satisfied with the customized file(s), you can then push the files to the network. Any subsequent changes to the files can be made on that same switch, and the new files imported and pushed to the network.

Editing a Switch File Set

Select the Switch File Set that you want to edit and click on the Edit icon . You can edit the **File Set Description** and **File Set Type**. When you are done editing, click on the **Apply** button. You can then assign the File Set to network devices.

Deleting a Switch File Set

Select the Switch File Set(s) that you want to delete and click on the Delete icon . Click **OK** at the confirmation prompt.

Settings

The Resource Manager Settings Screen is used to set the backup retention policy and the amount of space that must be available on the CMM before an upgrade is allowed.

Backup Retention Policy

These settings are used to specification of a maximum number of days and a minimum number of backups to keep per switch.

- **Minimum Backups** - The minimum number of backups you want to retain per switch. (Range = 1 - 365, Default = 365)
- **Backup Retention Period** - The maximum number of days that you want to retain those backups. (Range = 1 - 365, Default = 365)

If a backup for a switch is older than the maximum number of days, and the total number of backups is at least the minimum number specified, older backups will be deleted in accordance with the retention policy. The backup retention policy is applied when a new backup is successfully created.

BMF Upgrade Settings

- **Minimum Space** - The amount of space that must be available on the CMM before an upgrade is allowed. (Default = 4.5 MB)

22.0 SIP

Session Initiation Protocol (SIP) is an IETF-defined signaling protocol widely used for controlling communication sessions such as voice and video calls over Internet Protocol (IP). The protocol can be used for creating, modifying and terminating media sessions. SIP addresses the key challenge of real-time delivery and monitoring requirements for media streams from SIP devices. SIP Snooping prioritizes voice and video traffic over non-voice traffic. The OmniVista SIP Application automatically detects SIP data packets and enables you to configure SIP Profiles and apply QoS parameters for SIP packets; and monitor SIP traffic and create traps to alert you to SIP events. SIP Snooping:

- Identifies and marks the SIP and its corresponding media streams. Each media stream contains Real Time Protocol (RTP) and Real Time Control Protocol (RTCP) flows. Marking is done using the DSCP field in the IP header.
- Provides user configured QOS treatment for SIP/RTP/RTCP traffic flows based on its
- marking. Calculates QOS metric values of delay, jitter, round trip time, R factor and MOS values of media streams from its corresponding RTCP.

Note: The SIP snooping functions and the QOS actions require that the network paths used by the SIP signaling messages and the RTP/RTCP flows are the same and are “symmetric”. Therefore, MC-LAG, ECMP routing and VRRP topologies are not supported.

SIP

The screenshot displays the SIP configuration page in the Alcatel-Lucent Enterprise NMS. The page is titled "SIP" and includes a navigation menu on the left and a main content area with four tiles:

- Active Calls:** The Active Calls displays aggregated call data and detailed call data for any active calls on the network.
- Ended Calls:** The Ended Calls displays aggregated call data and detailed call data for any ended calls on the network.
- Sip Configuration:** Config One Touch, SIP, Global Param, Trusted Server, SOS, Threshold, TCP And UDP Port Profile.
- Setting:** Config SIP Call Setting.

SIP monitoring and configuration functions are accessed by clicking on one of the widgets on the Home Page or links on the left side of the screen.

- **Active Calls** - The Active Calls Screen displays call data for any active calls on the network.
- **Ended Calls** - The Ended Calls Screen displays call data for any ended calls on the network.
- **SIP Configuration** - SIP Configuration Screens are used to configure global SIP parameters and custom SIP Profiles. One or more of the following sub-profiles profiles can be included in a custom SIP Profile.
 - **One Touch Profile** - The One Touch Profile Screen is used to configure all parameters for SIP packets with a single command.
 - **SIP Profile** - SIP Profile Screens are used to configure custom SIP Profiles and assign the profiles to switches/ports in the network to specify how SIP traffic is handled.
 - **Global Param Profile** - Used to configure global SIP Profile parameters (e.g., DSCP marking, call thresholds) and enable/disable SIP Snooping.
 - **Trusted Servers Profile** - Used to configure the IP addresses of the Trusted Servers. If a Trusted Server is configured, only the calls initiated through those servers are supported. If no Trusted Servers are configured, all SIP based calls using any call server are supported.
 - **Threshold Profile** - Used to configure SIP Snooping threshold parameters (e.g., jitter, packet loss).
 - **SOS Profile** - Used to configure SOS call strings.
 - **TCP Port Profile** - Used to configure a TCP port(s) for SIP Snooping.
 - **UDP Port Profile** - Used to configure a UDP port(s) for SIP Snooping.
 - **Device View** - The Device View Screen displays SIP Profile configuration for any SIP-enabled switch in the network.
- **Settings** - The Settings Screen is used to enable/disable and configure data retention parameters for SIP data.

SIP Overview

Ever increasing applications and their need for network resources keep demand on networks high. Critical applications like real-time voice, video and mission critical data applications continue to grow, and bandwidth needs are growing at a faster pace than the network technologies that need to address them. Therefore it is essential to differentiate traffic, based on application, user and context, and provide applicable service levels for each. Voice and video traffic should be prioritized over non-voice traffic; and mission critical data traffic should be provided bandwidth guarantees for better performance. SIP is used for creating, modifying, and terminating media sessions; and applying QoS parameters to SIP traffic.

The SIP Snooping feature snoops voice quality metrics of media streams from their corresponding control packets and displays them to the user with knowledge of media reception quality in real time and helps to diagnose the problems on their quality. In addition, traps can be generated when voice/video/data quality parameters cross user configured thresholds.

Active Calls

The SIP Active Calls Screen is used to display Active Call Record data for selected SIP-enabled switches. To display Active Call Records, select an option from the drop-down menu (Use Switch Picker or Use Topology), then click on the **Select Devices** button to select the switches you want to view. The Active Call Records for the selected switches will be displayed in the table.

By default, the aggregated call records are displayed. The data is an aggregate of all Active Calls on SIP-enabled switches. You can also click on the **View Detailed Call** button at the top of the table to display detailed call records the selected switches.

Viewing Active Call Records

As described above, you can display aggregated or detailed call records in the table. You can also click on a switch(es) in the table to display a graphical representation of the call records.

Aggregated Records

Aggregated Records are call data for any active calls on the network. The data is an aggregate of all active calls on SIP-enabled switches.

- **Device** - The device name.
- **Start Time** - The call start date and time.
- **Calls Count** - The total number of calls processed for SIP Snooping.
- **RTCP Packet Count** - The total number of Real Time Control Protocol (RTCP) packets received by device.
- **RTP Packet Count** - The total number of Real Time Protocol (RTP) packet received by device.
- **Avg Pkt Loss** - The average number of SIP packet received by device.
- **Avg Jitter** - The average jitter, in milliseconds.
- **Avg RTD** - The average Round Trip Delay (RTD).
- **Avg RFactor** - The average RF Facto.
- **Avg MOS** - The average MOS.

Detailed Records

Detailed Records are detailed call data for any active calls on the network. The tab provides detailed data for each Active Call.

- **Device** - The device name.
- **Call ID** - The call ID.
- **Tag A** - The call tag for call direction A to B.
- **Tag B** - The call tag for call direction B to A.
- **IP Addr A Type** - The IP type for call direction A to B type (e.g., IPv4).
- **IP Addr A** - The IP address for call direction A to B.
- **IP Addr B Type** - The IP type for call direction B to A type (e.g., IPv4).
- **IP Address B** - The IP address for call direction B to A.
- **L4 Port A** - The call L4 port for call direction A to B.
- **L4 Port B** - The call L4 port for call direction B to A.
- **SIP Medial Type** - The SIP Media Type (e.g., Voice, Video)
- **Start Time** - The call start date and time.
- **RTP Count A** - The call Real Time Protocol (RTP) packet count for call direction A to B.
- **RTCP Type A** - The call Real Time Control Protocol (RTCP) packet count for call direction A to B.
- **Rule Name A** - The policy rule name for call direction A to B.
- **RTP Count B** - The call Real Time Protocol (RTP) packet count for call direction B to A.
- **RTCP Count B** - The call Real Time Protocol (RTP) packet count for call direction B to A.
- **Rule Name B** - The policy rule name for call direction B to A.
- **Jitter Violations A** - The call RTCP jitter violations (%) for call direction A to B.
- **Jitter Violations B** - The call RTCP jitter violations (%) for call direction B to A.
- **RTD Violation A** - The call round trip delay violations (%) for call direction A to B.

- **RTD Violation B** - The call round trip delay violations (%) for call direction B to A.
- **Packet Loss Violations A** - Call packet loss violations (%) for call direction A to B.
- **Packet Loss Violations B** - The call packet loss violations (%) for call direction B to A.
- **MOS Violations A** - The call MOS violations (%) for call direction A to B.
- **MOS Violations B** - The call MOS violations (%) for call direction B to A.
- **RF Factor Violations A** - The call RF Factor Violation (%) for call direction A to B.
- **RF Factor Violations B** - The call RF Factor Violation (%) for call direction B to A.
- **Jitter Max A** - The call maximum jitter for call direction A to B.
- **Jitter Min A** - The call minimum jitter for call direction A to B.
- **Jitter Avg A** - The call average jitter for call direction A to B.
- **Jitter Max B** - The call maximum jitter for call direction B to A.
- **Jitter Min B** - The call minimum jitter for call direction B to A.
- **Jitter Avg B** - The call average jitter for call direction B to A.
- **RTD Max A** - The call maximum round trip delay for direction A to B.
- **RTD Min A** - The call minimum round trip delay for direction A to B.
- **RTD Avg A** - The call average round trip delay for direction A to B.
- **RTD Max B** - The call maximum round trip delay for direction B to A.
- **RTD Min B** - The call minimum round trip delay for direction B to A.
- **RTD Avg B** - The call average round trip delay for direction B to A.
- **Pkt Loss Max A** - The call maximum packet loss (%) for call direction A to B.
- **Pkt Loss Min A** - The call minimum packet loss (%) for call direction A to B.
- **Pkt Loss Avg A** - The call average packet loss (%) for call direction A to B.
- **Pkt Loss Max B** - The call maximum packet loss (%) for call direction B to A.
- **Pkt Loss Min B** - The call minimum packet loss (%) for call direction B to A.
- **Pkt Loss Avg B** - The call average packet loss (%) for call direction B to A.
- **RF Factor Max A** - The call maximum RF Factor for call direction A to B.
- **RF Factor Min A** - The call minimum RF Factor for call direction A to B.
- **RF Factor Avg A** - The call average RF Factor for call direction A to B.
- **RF Factor Max B** - The call maximum RF Factor for call direction B to A.
- **RF Factor Min B** - The call minimum RF Factor for call direction B to A.
- **RF Factor Avg B** - The call average RF Factor for call direction B to A.
- **MOS Max A** - The call maximum MOS for call direction A to B.
- **MOS Min A** - The call minimum MOS for call direction A to B.
- **MOS Avg A** - The call average MOS for call direction A to B.
- **MOS Max B** - The call maximum MOS for call direction B to A.
- **MOS Min B** - The call minimum MOS for call direction B to A.
- **MOS Avg B** - The call average MOS for call direction B to A.

Graphical View

You can view a graphical representation of Active Call Records by selecting a switch or switches in the table. By default, the data for "Jitter" is displayed in bar chart format. However, you can select a different variable from the **Variable** drop-down menu; and also change the display to a pie chart by selecting the "Pie" radio button in the **Chart Type** area.

Ended Calls

The SIP Ended Calls Screen is used to display Ended Call Record data for selected SIP-enabled switches. To display Ended Call Records, select an option from the drop-down menu (Use Switch Picker or Use Topology), then click on the **Select Devices** button to select the switches you want to view. You can also configure a Start Time and End Time to only display records from a specific time period.

By default, the aggregated call records are displayed. The data is an aggregate of all Active Calls on SIP-enabled switches. You can also click on the **View Detailed Call** button at the top of the table to display detailed call records the selected switches.

Viewing Ended Call Records

As described above, you can display aggregated or detailed call records in the table. You can also click on a switch(es) in the table to display a graphical representation of the call records.

Aggregated Records

Aggregated Records display call data for any ended calls on the network. The data is an aggregate of all ended calls on SIP-enabled switches.

- **Device** - The device name.
- **Start Time** - The call start date and time.
- **End Time** - The call end date and time.
- **Calls Count** - The total number of calls processed for SIP Snooping.
- **RTCP Packet Count** - The total number of Real Time Control Protocol (RTCP) packets received by device.
- **RTP Packet Count** - The total number of Real Time Protocol (RTP) packet received by device.
- **Avg Pkt Loss** - The average number of SIP packet received by device.
- **Avg Jitter** - The average jitter, in milliseconds.
- **Avg RTD** - The average Round Trip Delay (RTD).
- **Avg RFactor** - The average RF Facto.
- **Avg MOS** - The average MOS.

Detailed Records

Detailed Records are detailed call data for any ended calls on the network. The tab provides detailed data for each ended call. The data is an aggregate of all ended calls.

- **Device** - The Device name.
- **Call ID** - The call ID.
- **Tag A** - The call tag for call direction A to B.
- **Tag B** - The call tag for call direction B to A.
- **IP Address A** - The IP address for call direction A to B.
- **IP Address B** - The IP address for call direction B to A.
- **Port A** - The call L4 port for call direction A to B.
- **Port B** - The call L4 port for call direction B to A.
- **Medial Type** - The SIP Media Type (e.g., Voice, Video)
- **Start Date** - The call start date and time.
- **End Date** - The call end date and time.
- **RTP Count A** - The call Real Time Protocol (RTP) packet count for call direction A to B.
- **RTCP Count A** - The call Real Time Control Protocol (RTCP) packet count for call direction

A to B.

- **Rule Name A** - The policy rule name for call direction A to B.
- **RTP Type B** - The RTP type for call direction B to A.
- **RTCP Count B** - The call Real Time Protocol (RTP) packet count for call direction B to A.
- **Rule Count B** - The call Real Time Protocol (RTP) packet count for call direction B to A.
- **Rule Name B** - The policy rule name for call direction B to A.
- **End Reason** - The end call reason.
- **Jitter Violation A** - The call RTCP jitter violations (%) for call direction A to B.
- **Jitter Violation B** - The call RTCP jitter violations (%) for call direction B to A.
- **RTD Violation A** - The call round trip delay violations (%) for call direction A to B.
- **RTD Violation B** - The call round trip delay violations (%) for call direction B to A.
- **Packet Loss Violation A** - The call packet loss violations (%) for call direction A to B.
- **Packet Loss Violation B** - The call packet loss violations (%) for call direction B to A.
- **MOS Violation A** - The call MOS violations (%) for call direction A to B.
- **MOS Violation B** - The call MOS violations (%) for call direction B to A.
- **RF Factor Violation A** - The call RF Factor Violation (%) for call direction A to B.
- **RF Factor Violation B** - The call RF Factor Violation (%) for call direction B to A.
- **Jitter Max A** - The call maximum jitter for call direction A to B.
- **Jitter Min A** - The call minimum jitter for call direction A to B.
- **Jitter Avg A** - The call average jitter for call direction A to B.
- **Jitter Max B** - The call maximum jitter for call direction B to A.
- **Jitter Min B** - The call minimum jitter for call direction B to A.
- **Jitter Avg B** - The call average jitter for call direction B to A.
- **RTD Max A** - The call maximum round trip delay for direction A to B.
- **RTD Min A** - The call minimum round trip delay for direction A to B.
- **RTD Avg A** - The call average round trip delay for direction A to B.
- **RTD Max B** - The call maximum round trip delay for direction B to A.
- **RTD Min B** - The call minimum round trip delay for direction B to A.
- **RTD Avg B** - The call average round trip delay for direction B to A.
- **Packet Loss Max A** - The call maximum packet loss (%) for call direction A to B.
- **Packet Loss Min A** - The call minimum packet loss (%) for call direction A to B.
- **Packet Loss Avg A** - The call average packet loss (%) for call direction A to B.
- **Packet Loss Max B** - The call maximum packet loss (%) for call direction B to A.
- **Packet Loss Min B** - The call minimum packet loss (%) for call direction B to A.
- **Packet Loss Avg B** - The call average packet loss (%) for call direction B to A.
- **RF Factor Max A** - The call maximum RF Factor for call direction A to B.
- **RF Factor Min A** - The call minimum RF Factor for call direction A to B.
- **RF Factor Avg A** - The call average RF Factor for call direction A to B.
- **RF Factor Max B** - The call maximum RF Factor for call direction B to A.
- **RF Factor Min B** - The call minimum RF Factor for call direction B to A.
- **RF Factor Avg B** - The call average RF Factor for call direction B to A.
- **MOS Max A** - The call maximum MOS for call direction A to B.
- **MOS Min A** - The call minimum MOS for call direction A to B.
- **MOS Avg A** - The call average MOS for call direction A to B.
- **MOS Max B** - The call maximum MOS for call direction B to A.
- **MOS Min B** - The call minimum MOS for call direction B to A.
- **MOS Avg B** - The call average MOS for call direction B to A.

Graphical View

You can view a graphical representation of Active Call Records by selecting a switch or switches in the table. By default, the data for "Jitter" is displayed in bar chart format. However, you can select a different variable from the **Variable** drop-down menu; and also change the display to a pie chart by selecting the "Pie" radio button in the **Chart Type** area.

One Touch Profile

The SIP One Touch Profile Screen displays all configured SIP One Touch Profiles, and is used to create, edit, delete, and apply One Touch Profiles. A Custom SIP Profile can be created and applied to switches on the network using the SIP Profile Tab. However, a SIP One Touch Profile is an easy way to automatically apply a default SIP Profile to traffic on SIP-enabled switches.

SIP One Touch Profile Parameters

When a SIP One Touch Profile is created and applied to a switch on the network, the following SIP configuration is applied to that switch:

Edge Devices

- The SIP Snooping Status is set to "Enabled" and the SOS Call Number values configured based on Media Type:
 - Voice - SOS Call Number 1 Field for the SOS Call Number.
 - Video - SOS Call Number 2 Field for the SOS Call Number.
 - Other - SOS Call Number 2 Field for the SOS Call Number.
- The Port Mode for all ports on the switch is set to "Automatic".
- The Port Status for all ports on the switch is set to "Enabled".
- A One Touch Policy Rule is created on the switch.

Non-Edge Devices

- The SIP Snooping Status is set to "Enabled" and the SOS Call Number values configured based on Media Type:
 - Voice - SOS Call Number 1 Field for the SOS Call Number.
 - Video - SOS Call Number 2 Field for the SOS Call Number.
 - Other - SOS Call Number 2 Field for the SOS Call Number.
- A One Touch Policy Rule is created on the switch.
- The user will have to manually set the Port Mode to Force-Edge/Force-Non-Edge from the CLI or by creating and assigning a custom SIP Profile using the SIP Profile Screen.

One Touch Policy Rule

The One Touch Policy Rule for different media types is detailed below.

- Voice
 - Policy Condition - sip audio Policy Action - dscp 46
 - Policy Rule - OneTouchSIPRule\$Voice condition OneTouchSIPCondition\$Voice action OneTouchSIPAction\$Voice

- Video
 - Policy Condition - sip video Policy Action - dscp 34
 - Policy Rule - OneTouchSIPRule\$Video condition OneTouchSIPCondition\$Video action OneTouchSIPAction\$Video
- Other
 - Policy Condition - sip other
 - Policy Action - dscp 24
 - Policy Rule - OneTouchSIPRule\$Other condition OneTouchSIPCondition\$Other action OneTouchSIPAction\$Other

SIP One Touch Policy Rule Precedence

PolicyView enables you to define the precedence of policies created in PolicyView. A policy rule's precedence determines which policy will take effect in the rare case of a conflict.


- One Touch SIP Voice Policy Precedence is fixed at 50000.
- One Touch SIP Video Policy Precedence is fixed at 44000.
- One Touch SIP Other Policy Precedence is fixed at 44001.

Note the following for SIP Policies:

- Precedence for two policies can be the same.
- SIP Voice Precedence should be higher than Video/Other precedence.
- SIP Voice Precedence can overlap in One Touch Voice range.

Note See the PolicyView Help for more information on Policy Precedence and Conflicts.


Creating a SIP One Touch Profile

SIP One Touch Profiles must be unique. You cannot create two profiles with the same name; and you cannot create duplicate profiles. Two profiles are considered duplicate if they have the same SIP Media Type and SOS Call Number. To create a SIP One Touch Profile, click on the Create icon . Click on one or more Media Types (Video, Voice, Other) and enter an SOS Call Number. When you are finished, click on the **Create** button. The profile will be created and stored in OmniVista.


After creating a profile, select the profile in the One Touch Profile List and click on the Apply to Devices button to apply the profile to specific network switches. The profile will be assigned to all ports on the selected switches.

Note: The SIP snooping features allow the detection of emergency calls based on the “to” URI in the invite message. The SOS string must be the exact URI to be matched in the “to” URI; regular expressions are not supported.

Editing a SIP One Touch Profile

You can edit the SOS Call Numbers for a profile. Select a profile in the One Touch Profile List and click on the Edit icon . Edit the SOS Call Number(s) and click on the **Apply** button. The edited profile will be applied to any assigned switches/ports. Note that you cannot edit the profile name.

Deleting a SIP One Touch Profile

Select a profile in the One Touch Profile List, click on the Delete icon  then click **OK** at the Confirmation Prompt. The profile will be deleted from the One Touch Profile List and from all assigned switches/ports.

Removing a SIP One Touch Profile

To remove an assigned One Touch Profile from specific switches/ports, select the profile in the SIP List and click on the **Apply to Devices** button. Select an option from the drop-down menu (Use Switch Picker or Use Topology) and click on the **Add/Remove Devices** button to select the switch(es) from which you want to remove the profile. Move the switch(es) from the Selected column and click on the **Apply** button.

Note: Removing a SIP One Touch Profile from a switch does not change the global SIP Snooping Status, SOS call number, or port parameters.

Applying a SIP One Touch Profile

Select a profile in the One Touch Profile List and click on the **Apply to Devices** button. Select an option from the drop-down menu (Use Switch Picker or Use Topology) and click on the **Add/Remove Devices** button to select the switch(es) to which you want to apply the profile. Select a **Device Type** from the drop-down menu (Edge or Non Edge) and click on the **Apply** button. The profile will be assigned to all ports on the selected switches.

Note: You can only apply one (1) One Touch Policy to a switch. Only switches without an applied One Touch Policy will be available for selection. If you want to apply a new One Touch Policy to a switch, you must first remove the existing policy.

Note: SIP Snooping is not supported in a Multi-Chassis configuration. If Multi-Chassis is configured on a 9000E Series device, that device will not be visible in the device list. If a device is configured in a Multi-Chassis configuration after opening the SIP application, assigning the device will result in an error Message. If Multi-Chassis is not configured and a device is not visible in the device list, the device must be polled so that its status is updated and it will appear in the list.

Viewing SIP One Touch Profiles

The One Touch Profile List displays information for all configured One Touch Profiles. Click on a profile in the list to display profile details, including the switches to which the profile was assigned, if applicable.

- **Profile Name** - The user-configured name for the profile
- **Video SOS Call** - The SOS Call Number for Video, if configured for the profile.
- **Voice SOS Call** - The SOS Call Number for Voice, if configured for the profile.
- **Other SOS Call** - The SOS Call Number for other media types, if configured for the profile.

SIP Profile

The SIP Profile Screen displays all configured SIP Profiles and is used create, edit, and delete SIP Profiles and apply profiles to switches/ports in the network. A SIP Profile is basically a "Master" Profile made up of "sub-profile" parameters configured using Global Params, Trusted Servers, Threshold, Threshold, TCP Port, UDP Port, and SOS Profile Screens.

Creating a SIP Profile


To create a SIP Profile, click on the Create icon **+** and enter a **SIP Profile Name**. Click on the "Add More Profiles" and select any sub-profile types you want to include. Select a sub-profile from a drop-down. If necessary, click on the "Add New" link to go to a sub-profile page and create a new profile, then select the profile from the drop-down list. When you are finished, click on the **Create** button. The profile will be created and stored in OmniVista.

- **Global Params Profile** - Used to configure global SIP Profile parameters (e.g., DSCP marking, call thresholds) and enable/disable SIP Snooping.
- **Trusted Servers Profile** - Used to configure the IP addresses of the Trusted Servers. If a Trusted Server is configured, only the calls initiated through those servers are supported. If no Trusted Servers are configured, all SIP based calls using any call server are supported.
- **Threshold Profile** - Used to configure SIP Snooping threshold parameters (e.g., jitter, packet loss).
- **SOS Profile** - Used to configure SOS call strings.
- **TCP Port Profile** - Used to configure a TCP port(s) for SIP Snooping.
- **UDP Port Profile** - Used to configure a UDP port(s) for SIP Snooping.


After creating a profile, select the profile in the SIP Profile List and click on the Apply to Devices button to apply the profile to specific network switches/ports.

Note: You can only have one of each sub-profile type in a SIP Profile. In other words you cannot have two (2) different Trusted Server Profiles or two (2) different SOS Profiles.

Editing a SIP Profile

You can add/remove sub-profiles to/from a SIP Profile. Select a profile in the SIP Profile List and click on the Edit icon . You can select different sub-profiles, click on the Add More Profiles Link to add additional sub-profiles to the SIP Profile, or click on the "Remove" link next to a sub-profile to remove it from the SIP Profile. When you are finished, click on the **Apply** button. The profile will be updated and re-applied to any assigned switches/ports.

Deleting a SIP Profile

Select a profile in the SIP Profile List, click on the Delete icon  then click **Yes** at the Confirmation Prompt. The profile will be deleted from the SIP Profile List and from all assigned switches/ports.

To remove an assigned SIP Profile from specific switches/ports, select the profile in the SIP List and click on the Apply to Devices button. Select an option from the drop-down menu (Use Switch Picker or Use Topology) and click on the Add/Remove Devices button to select the switch(es) from which you want to remove the profile. Move the switch(es) from the Selected column and click on the Apply button.

To remove the profile from specific ports, click on a switch and click on the Add/Remove Ports button to select the ports from which you want to remove the profile. Move the port(s) from the Selected column and click on the **Apply** button.

Applying a SIP Profile

Select a profile in the SIP Profile List and click on the Apply to Devices button. Select an option from the drop-down menu (Use Switch Picker or Use Topology) and click on the Add/Remove Devices button to select the switch(es) to which you want to apply the profile. Click on a switch and click on the Add/Remove Ports button to select the ports to which you want to apply the profile. (You can also click on the "Add Port" link to bring up a list of ports for selection.)

By default, the Port Mode is set to "Automatic". To change the port mode, click on the "Device Config" Link for the switch to bring up the Device Configuration window and select a different Port Mode:

- **Automatic** - The port Edge/Non-Edge mode is derived from the switch based on LLDP received on the port.
- **Force Edge** - Media TCAM entries are created for dialogs that transverse the port.
- **Force Non-Edge** - Media TCAM entries are not created for dialogs that transverse the port.

Note that a port on a device will be considered as an Edge Port if it is" not connected through a link to any other port.

- It is connected through a link to another port on which LLDP is disabled.
- It has LLDP enabled and is connected through a link to another port on which LLDP is enabled and the remote capability advertised by that port is "None".

A port on a device will be considered as a Non-Edge Port if:

- It has LLDP enabled and is connected through a link to another port on which LLDP is enabled and the remote capability advertised by that port is either "Bridge" or "Router".

When you are finished, click on the **Apply** button. The profile will be applied to the selected switch ports.

Note: You can only apply one (1) SIP Policy to a switch. Only switches without an applied One Touch Policy will be available for selection. If you want to apply a new SIP Policy to a switch, you must first remove the existing policy.

Note: SIP Snooping is not supported in a Multi-Chassis configuration. If Multi-Chassis is configured on a 9000E Series device, that device will not be visible in the device list. If a device is configured in a Multi-Chassis configuration after opening the SIP application, assigning the device will result in an error Message. If Multi-Chassis is not configured and a device is not visible in the device list, the device must be polled so that its status is updated and it will appear in the list.

Viewing SIP Profiles

The SIP Profile List displays information for all configured SIP Profiles. Click on a profile in the list to display profile details, including the switches to which the profile was assigned, if applicable.


- **SIP Profile Name** - The user-configured name for the SIP Profile.
- **SIP Profile Status**
 - **In Sync** - The sub-profiles contained in the SIP Profile have not changed since the profile was created.
 - **Out of Sync** - A sub-profile contained in the SIP Profile as been edited since it was initially included in the SIP Profile. Any switches/ports to which the profile was initially applied will retain the original SIP Profile configuration until the profile is re-applied to the switches/ports.
 - **Unassigned** - The SIP profiled has not yet been applied to switches/ports.
- **Global Params Profile** - The name of the Global Parameters Profile contained in the SIP Profile. Global Parameters include SIP Snooping Enable/Disable, DSCP Marking, and Call Thresholds.
- **Trusted Servers Profile** - The name of the Trusted Servers Profile contained in the SIP Profile. A Trusted Server Profile contains IP addresses of Trusted Servers. If a Trusted Server is configured, only the calls initiated through those servers are supported. If no Trusted Servers are configured, all SIP based calls using any call server are supported.
- **UDP Ports Profile** - The name of the UDP Ports Profile contained in the SIP Profile. A UDP Ports Profile contains TCP Ports configured for SIP Snooping.
- **TCP Ports Profile** - The name of the TCP Ports Profile contained in the SIP Profile. A TCP Ports Profile contains TCP Ports configured for SIP Snooping.
- **Threshold Profile** - The name of the Threshold Profile contained in the SIP Profile. A Threshold Profile contains SIP Snooping threshold parameters (e.g., Jitter, Packet Loss).
- **SOS Profile** - The name of the SOS Profile contained in the SIP Profile. An SOS Profile contains a list of SOS call strings.

Note: The SIP Profile Name only displays if a SIP Profile exists on the switch that completely matches the applied profile configuration; otherwise the SIP Profile Name field is blank.

Global Params Profile


The SIP Global Params Profile Screen displays all configured SIP Global Parameters Profiles and is used create, edit, and delete Global Parameter Profiles. A Global Parameters Profile can then be included in a SIP Profile and assigned to switches/ports in the network.

Creating a Global Parameters Profile

Global Parameter Profiles must be unique. You cannot create two profiles with the same name; and you cannot create duplicate profiles. Two Global Parameter Profiles are considered duplicate if they have the same DSCP Number, SOS Call DSCP Number, and Threshold Number of Calls values. To create a Global Parameters Profile, click on the Create icon  and complete the fields as described below. When you are finished, click on the **Create** button.


- **Profile Name** - User-configured profile name (up to 32 characters).
- **DSCP Number**- The SIP Snooping DSCP number. By default, the packet gets its priority as normal packet (Range = 0 to 63). If the DSCP Number field is set to "0", the value is set as "NA" on switch.
- **SOS Call DSCP No.** - The SIP Snooping SOS Call DSCP No. (Range = 0 to 63)
- **Threshold No. of Calls** - The Number of call records that can be stored in flash (Range = 50 to 500)
- **Reserved Hardware Resource** - Reserved hardware resources required to program ACLs for media entries. Each value is a multiple of the default reserved hardware resources.
- **SIP CPU Rate Limit** - The rate limit of SIP PDUs trapping toward CPU (not applicable to SIP PDUs going towards network port).
- **SIP Snooping Status** - Use the drop-down menu to Enable/Disable SIP Snooping.
- **Clear Stats** - If set to "Yes", when the profile is assigned, existing SIP Statistics are cleared.

Editing a Global Parameters Profile

Select a profile in the Global Params Profile List and click on the Edit icon . Edit the fields as described above and click on the **Apply** button (note that you cannot edit the profile name). If you edit a Global Parameters Profile that is part of a SIP Profile, the SIP Profile will be labeled as "Out of Sync" on the SIP Profile Screen. Any device(s)/port(s) that the now-edited SIP Profile was applied to will retain its current SIP configuration until the SIP Profile is re-applied to those devices.

To apply the edited SIP Profile, go to the SIP Profiles Screen, select the SIP Profile that includes the edited Global Parameters Profile, and click on the **Apply to Devices** button. Select the device(s)/port(s) to which you want to re-apply the updated SIP Profile, and click on the **Apply** button.

Deleting a Global Parameters Profile

Select a profile in the Global Params Profile List, click on the Delete icon  then click **Yes** at the Confirmation Prompt. Note that you cannot delete a Global Parameters Profile that has been applied to a switch as part of a SIP Profile. You must first remove the SIP Profile from the switch.

Viewing Global Parameters Profiles

The Global Params Profile List displays information for all configured Global Parameter Profiles.

- **Profile Name** - User-configured profile name.
- **DSCP Number**- The SIP Snooping DSCP number. By default, the packet gets its priority as normal packet (Range = 0 to 63). If the DSCP Number field is set to "0", the value is set as "NA" on switch.

- **SOS Call DSCP No.** - The SIP Snooping SOS Call DSCP No. (Range = 0 to 63)
- **Threshold No. of Calls** - The Number of call records that can be stored in flash (Range = 50 to 500)
- **Clear Stats** - If set to "Yes", when the profile is assigned, existing SIP Statistics are cleared.
- **Reserved Hardware Resource** - Reserved hardware resources required to program ACLs for media entries. Each value is a multiple of the default reserved hardware resources.
- **SIP CPU Rate Limit** - The rate limit of SIP PDUs trapping toward CPU (not applicable to SIP PDUs going towards network port).

Trusted Servers Profile


The SIP Trusted Servers Profile Screen displays all configured Trusted Servers Profiles and is used create, edit, and delete Trusted Servers Profiles. If a Trusted Server is configured, only calls initiated through those servers are supported. If no Trusted Servers are configured, all SIP based calls using any call server are supported. You can configure up to eight (8) IP Addresses. A Trusted Server Profile can be included in a SIP Profile and assigned to switches/ports in the network.


Creating a Trusted Servers Profile

When creating a Trusted Servers Profile, the Profile Name, and at least one IP Address are required. Profiles must be unique. You cannot create two (2) profiles with the same name; and you cannot create duplicate profiles. Two Trusted Server Profiles are considered duplicate if they have the same IP Addresses in the same IP Address field. For example, if Profile 1 was created with IP Address 1 and IP Address 2 specified as


1.1.1.1 and 2.2.2.2; and Profile 2 was created with IP Address 1 and IP Address 2 specified as 1.1.1.1 and 2.2.2.2; and Profile 3 was created with IP Address 2 and IP Address 3 specified as 1.1.1.1 and 2.2.2.2; Profile 1 and Profile 2 would be duplicate profiles, but Profile 1 and Profile 3 would not.

Note: When you apply a new Trusted Servers Profile as part of a SIP Profile it completely removes the previous Trusted Servers configuration on the switch and configures the Trusted Server IP Addresses provided in the new profile.

To create a Trusted Servers Profile, click on the Create icon  and complete the fields as described below. When you are finished, click on the **Create** button.

- **Profile Name** - User-configured profile name (up to 32 characters).
- **IP Address List** - Enter a Trusted Server IP address and click on the Add icon . Repeat to add additional servers. You can configure up to eight (8) IP Addresses. At least one IP Address is required for the profile. "0.0.0.0" is considered as an invalid Trusted Server IP Address.

Editing a Trusted Servers Profile

Select a profile in the Trusted Servers Profile List and click on the Edit icon . Edit the fields as described above and click on the **Apply** button (note that you cannot edit the profile name). If you edit a Trusted Servers Profile that is part of a SIP Profile, the SIP Profile will be labeled as "Out of Sync" on the SIP Profile Screen. Any device(s)/port(s) that the now-edited SIP Profile was applied to will retain its current SIP configuration until the SIP Profile is re-applied to those devices.

To apply the edited SIP Profile, go to the SIP Profiles Screen, select the SIP Profile that includes the edited Trusted Servers Profile, and click on the **Apply to Devices** button. Select the device(s)/port(s) to which you want to re-apply the updated SIP Profile, and click on the **Apply** button.

Deleting a Trusted Servers Profile

Select a profile in the Trusted Servers Profile List, click on the Delete icon  then click **Yes** at the

Confirmation Prompt. Note that you cannot delete a Trusted Servers Profile that has been applied to a switch as part of a SIP Profile. You must first remove the SIP Profile from the switch.

Viewing Trusted Servers Profiles

The Trusted Servers List displays information for all configured Trusted Servers Profiles.

- **Profile Name** - User-configured profile name.
- **IP Address List** - IP addresses of Trusted Servers included in the profile.

Threshold Profile

The SIP Threshold Profile Screen displays all configured Threshold Profiles and is used create, edit, and delete Threshold Profiles for SIP Snooping (e.g., jitter, packet loss). A Threshold Profile can be included in a SIP Profile and assigned to switches/ports in the network. (e.g., jitter, packet loss).

Creating a Threshold Profile


When creating a Threshold Profile, the Profile Name, at least one (1) medium (e.g., Audio, Video Other), and at least one of the performance parameter fields (e.g., Jitter, Packet Loss) are required. Profiles must be unique. You cannot create two profiles with the same name; and you cannot create duplicate profiles. Two Threshold Profiles are considered duplicate if they have the same values in the performance parameter fields (e.g., Jitter, Packet Loss). Follow the steps below to create a Threshold Profile.

To create a Threshold Profile, click on the Create icon  and complete the fields for one or more medium as described below. When you are finished, click on the **Create** button.

- **Profile Name** - User-configured profile name (up to 32 characters).
- **Jitter** - The Jitter Threshold, in milliseconds (Range = 0 to 300, Defaults = Audio - 50, Video - 100, Other - 100).
- **Packet Loss** - The Packet Loss Threshold, in % (Range = 0 to 99, Defaults = Audio - 10, Video - 20, Other - 20).
- **Round Trip Delay** - The Round Trip Delay Threshold, in milliseconds (Range = 0 to 500, Defaults = Audio - 80, Video - 250, Other - 250).
- **R Factor** - The R-Factor Threshold, in milliseconds (Range = 0 - 100, Defaults = Audio - 70, Video - 80, Other - 80).
- **MOS** - The MOS Value Threshold (Range = 0 - 5, Defaults = Audio - 3.6, Video - 3.0, Other - 3.0). Note that the previous MOS range was 0 - 50. The current range of 0 - 5 represents 1/10th of the previous values.

Note: Applying a new Threshold Profile will modify the threshold values for the specified performance parameter fields.


Editing a Threshold Profile

Select a profile in the Threshold Profile List and click on the Edit icon . Edit the fields as described above and click on the **Apply** button (note that you cannot edit the profile name). If you edit a Threshold Profile that is part of a SIP Profile, the SIP Profile will be labeled as "Out of Sync" on the SIP Profile Screen. Any device(s)/port(s) that the now-edited SIP Profile was applied to will retain its current SIP configuration until the SIP Profile is re-applied to those devices.

To apply the edited SIP Profile, go to the SIP Profiles Screen, select the SIP Profile that includes the edited Threshold Profile, and click on the **Apply to Devices** button. Select the device(s)/port(s) to which you want

to re-apply the updated SIP Profile, and click on the **Apply** button.

Deleting a Threshold Profile

Select a profile in the Threshold Profile List, click on the Delete icon  then click **Yes** at the Confirmation Prompt. Note that you cannot delete a Threshold Profile that has been applied to a switch as part of a SIP Profile. You must first remove the SIP Profile from the switch.

Viewing Threshold Profiles

The Threshold Profile List displays information for all configured Threshold Profiles.

- **Profile Name** - User-configured profile name.
- **Jitter** - The Jitter Threshold, in milliseconds (Range = 0 to 300, Defaults = Audio - 50, Video - 100, Other - 100).
- **Packet Loss** - The Packet Loss Threshold, in % (Range = 0 to 99, Defaults = Audio - 10, Video - 20, Other - 20).
- **Round Trip Delay** - The Round Trip Delay Threshold, in milliseconds (Range = 0 to 500, Defaults = Audio - 80, Video - 250, Other - 250).
- **R Factor** - The R-Factor Threshold, in milliseconds (Range = 0 - 100, Defaults = Audio - 70, Video - 80, Other - 80).
- **MOS** - The MOS Value Threshold (Range = 0 - 5, Defaults = Audio - 3.6, Video - 3.0, Other - 3.0). Note that the previous MOS range was 0 - 50. The current range of 0 - 5 represents 1/10th of the previous values.

SOS Profile


The SIP SOS Profile Screen displays all configured SOS Profiles and is used create, edit, and delete SOS Profiles for SIP Snooping. A SOS Profile can be included in a SIP Profile and assigned to switches/ports in the network. (e.g., jitter, packet loss).

Creating an SOS Profile

When creating an SOS Profile, the Profile Name, and at least one SOS Call Number are required. Profiles must be unique. You cannot create two profiles with the same name; and you cannot create duplicate profiles. Two SOS Profiles are considered duplicate if they have the same SOS Call Number in the same SOS Call Number field. For example, if Profile 1 was created with SOS Call Number 1 and SOS Call Number 2 specified as abcd and 1234; and Profile 2 was created with SOS Call Number 1 and SOS Call Number 2 specified as abcd and 1234; and Profile 3 was created with SOS Call Number 2 and SOS Call Number 3 specified as abcd and 1234; Profile 1 and Profile 2 would be duplicate profiles, but Profile 1 and Profile 3 would not.

Note: When you apply a new SOS Profile it completely removes the previous SOS Call Number configuration on the switch and configures the SOS Call Numbers provided in the new profile.




To create an SOS Profile, click on the Create icon  and complete the fields for one or more medium as described below. When you are finished, click on the **Create** button.

- **Profile Name** - User-configured profile name (up to 32 characters).
- **SOS Call Number List** - Enter a call string for the profile and click on the Add icon . Repeat to add additional numbers. The SIP Snooping features allow the detection of emergency calls based on the “to” URI in the invite message. You can configure up to Configuration allows up to 4 SOS call strings. The string must be the exact URI to be matched in the “to” URI; regular expression is not supported.

By default, no SOS number is configured for SIP Snooping. (Allowed characters are a-z, A-Z, 0-9, @.)

Note: If the SOS Call Number field is left blank, when the profile is assigned to a switch, OmniVista will erase that SOS Call Number on the switch, if it existed previously.


Editing an SOS Profile

Select a profile in the SOS Profile List and click on the Edit icon . Edit the fields as described above. Click on Add icon  and enter a number to add an additional number. Click on the Delete icon  to remove an existing number. When you are finished, click on the **Apply** button (note that you cannot edit the profile name).

If you edit an SOS Profile that is part of a SIP Profile, the SIP Profile will be labeled as "Out of Sync" on the SIP Profile Screen. Any device(s)/port(s) that the now-edited SIP Profile was applied to will retain its current SIP configuration until the SIP Profile is re-applied to those devices.

To apply the edited SIP Profile, go to the SIP Profiles Screen, select the SIP Profile that includes the edited SOS Profile, and click on the **Apply to Devices** button. Select the device(s)/port(s) to which you want to re- apply the updated SIP Profile, and click on the **Apply** button.

Deleting an SOS Profile

Select a profile in the SOS Profile List, click on the Delete icon  then click **Yes** at the Confirmation Prompt. Note that you cannot delete an SOS Profile that has been applied to a switch as part of a SIP Profile. You must first remove the SIP Profile from the switch.

Viewing SOS Profiles

The SOS Profile List displays information for all configured SOS Profiles.

- **Profile Name** - User-configured profile name.
- **SOS Call Number List** - The SOS number(s) configured for the profile.

TCP Port Profile

The SIP TCP Port Profile Screen displays all configured TCP Port Profiles and is used create, edit, and delete TCP Port Profiles. A TCP Port Profile can be included in a SIP Profile and assigned to switches/ports in the network. You can configure up to eight (8) TCP Ports on SIP Snooping devices.


Creating a TCP Port Profile

When creating a TCP Port Profile, the Profile Name, and at least one TCP Port are required. Profiles must be unique. You cannot create two profiles with the same name; and you cannot create duplicate profiles. Two TCP Port Profiles are considered duplicate if they have the same TCP Port Number in the same TCP Port field. For example, if Profile 1 was created with TCP Port 1 and TCP Port 2 specified as 30 and 40; and Profile 2 was created with TCP Port 1 and TCP Port 2 specified as 30 and 40; and Profile 3 was created with TCP Port 2 and TCP Port 3 specified as 30 and 40; Profile 1 and Profile 2 would be duplicates, but Profile 1 and Profile 3 would not.




Note: When you apply a new TCP Port Profile it completely removes the previous TCP Ports configuration on the switch and configures the TCP Ports provided in the new profile.

To create a TCP Port Profile, click on the Create icon  and complete the fields for one or more medium as

described below. When you are finished, click on the **Create** button.

- **Profile Name** - User-configured profile name (up to 32 characters).
- **TCP Port List** - Enter a TCP Port Number for the profile and click on the Add icon . Repeat to add additional ports. You can configure up to eight (8) TCP Ports on SIP Snooping devices (Port Range 0 - 65535).


Editing a TCP Port Profile

Select a profile in the TCP Port Profile List and click on the Edit icon . Edit the fields as described above. Click on Add icon  and enter a port to add an additional port. Click on the Delete icon  to remove an existing port. When you are finished, click on the **Apply** button (note that you cannot edit the profile name).

If you edit a TCP Port Profile that is part of a SIP Profile, the SIP Profile will be labeled as "Out of Sync" on the SIP Profile Screen. Any device(s)/port(s) that the now-edited SIP Profile was applied to will retain its current SIP configuration until the SIP Profile is re-applied to those devices.

To apply the edited SIP Profile, go to the SIP Profiles Screen, select the SIP Profile that includes the edited TCP Port Profile, and click on the **Apply to Devices** button. Select the device(s)/port(s) to which you want to re-apply the updated SIP Profile, and click on the **Apply** button.

Deleting a TCP Port Profile

Select a profile in the TCP Port Profile List, click on the Delete icon  then click **Yes** at the Confirmation Prompt. Note that you cannot delete a TCP Port Profile that has been applied to a switch as part of a SIP Profile. You must first remove the SIP Profile from the switch.

Viewing TCP Port Profiles

The TCP Port Profile List displays information for all configured TCP Port Profiles.

- **Profile Name** - User-configured profile name.
- **TCP Port List** - The TCP Port(s) configured for the profile.

UDP Port Profile

The SIP UDP Port Profile Screen displays all configured UDP Port Profiles and is used create, edit, and delete UDP Port Profiles. A UDP Port Profile can be included in a SIP Profile and assigned to switches/ports in the network. You can configure up to eight (8) UDP Ports on SIP Snooping devices.


Creating a UDP Port Profile

When creating a UDP Port Profile, the Profile Name, and at least one UDP Port are required. Profiles must be unique. You cannot create two profiles with the same name; and you cannot create duplicate profiles. Two UDP Port Profiles are considered duplicate if they have the same UDP Port Numbers in the same UDP Port field. For example, if Profile 1 was created with UDP Port 1 and UDP Port 2 specified as 30 and 40; and Profile 2 was created with UDP Port 1 and UDP Port 2 specified as 30 and 40; and Profile 3 was created with UDP Port 2 and UDP Port 3 specified as 30 and 40; Profile 1 and Profile 2 would be duplicates, but Profile 1 and Profile 3 would not.




Note: When you apply a new UDP Port Profile it completely removes the previous UDP Ports configuration on the switch and configures the UDP Ports provided in the new profile.

To create a UDP Port Profile, click on the Create icon  and complete the fields for one or more medium as

described below. When you are finished, click on the **Create** button.

- **Profile Name** - User-configured profile name (up to 32 characters).
- **UDP Port List** - Enter a UDP Port Number for the profile and click on the Add icon . Repeat to add additional ports. You can configure up to eight (8) UDP Ports on SIP Snooping devices (Port Range 0 - 65535).


Editing a UDP Port Profile

Select a profile in the UDP Port Profile List and click on the Edit icon . Edit the fields as described above. Click on Add icon  and enter a port to add an additional port. Click on the Delete icon  to remove an existing port. When you are finished, click on the **Apply** button (note that you cannot edit the profile name).

If you edit an UDP Port Profile that is part of a SIP Profile, the SIP Profile will be labeled as "Out of Sync" on the SIP Profile Screen. Any device(s)/port(s) that the now-edited SIP Profile was applied to will retain its current SIP configuration until the SIP Profile is re-applied to those devices.

To apply the edited SIP Profile, go to the SIP Profiles Screen, select the SIP Profile that includes the edited UDP Port Profile, and click on the **Apply to Devices** button. Select the device(s)/port(s) to which you want to re-apply the updated SIP Profile, and click on the **Apply** button.

Deleting a UDP Port Profile

Select a profile in the UDP Port Profile List, click on the Delete icon  then click Yes at the Confirmation Prompt. Note that you cannot delete a UDP Port Profile that has been applied to a switch as part of a SIP Profile. You must first remove the SIP Profile from the switch.

Viewing UDP Port Profiles

The UDP Port Profile List displays information for all configured UDP Port Profiles.

- **Profile Name** - User-configured profile name.
- **UDP Port List** - The UDP Port(s) configured for the profile.

Device View

The SIP Device View Screen displays the SIP Profile configuration and SIP statistics for any SIP-enabled switch in the network. Select an option from the drop-down menu (Use Switch Picker/Use Topology) and click on the **Select Device** button to select a switch. Click on a profile type to view configuration information. The following SIP configuration information is displayed for the selected switch.

- SIP Profile
- TCP Port Profile UDP
- Port Profile
- Global Param Profile
- Trusted Server
- Profile Threshold
- Profile SOS Profile
- SIP Statistics

SIP Profile

- **SIP Profile Name** - The user-configured name for the SIP Profile.
- **SIP Profile Status**
 - **In Sync** - The sub-profiles contained in the SIP Profile have not changed since the profile was created.
 - **Out of Sync** - A sub-profile contained in the SIP Profile as been edited since it was initially included in the SIP Profile. Any switches/ports to which the profile was initially applied will retain the original SIP Profile configuration until the profile is re-applied to the switches/ports.
 - **Unassigned** - The SIP profiled has not yet been applied to switches/ports.
- **Global Params Profile** - The name of the Global Parameters Profile contained in the SIP Profile. Global Parameters include SIP Snooping Enable/Disable, DSCP Marking, and Call Thresholds.
- **Trusted Servers Profile** - The name of the Trusted Servers Profile contained in the SIP Profile. A Trusted Server Profile contains IP addresses of Trusted Servers. If a Trusted Server is configured, only the calls initiated through those servers are supported. If no Trusted Servers are configured, all SIP based calls using any call server are supported.
- **UDP Ports Profile** - The name of the UDP Ports Profile contained in the SIP Profile. A UDP Ports Profile contains TCP Ports configured for SIP Snooping.
- **TCP Ports Profile** - The name of the TCP Ports Profile contained in the SIP Profile. A TCP Ports Profile contains TCP Ports configured for SIP Snooping.
- **Threshold Profile** - The name of the Threshold Profile contained in the SIP Profile. A Threshold Profile contains SIP Snooping threshold parameters (e.g., Jitter, Packet Loss).
- **SOS Profile** - The name of the SOS Profile contained in the SIP Profile. An SOS Profile contains a list of SOS call strings.

TCP Port Profile

- **Profile Name** - User-configured TCP Port Profile name.
- **TCP Port List** - The TCP Port(s) configured for the profile.

UDP Port Profile

- **Profile Name** - User-configured UDP Port Profile name (up to 32 characters).
- **UDP Port List** - The UDP Port(s) configured for the profile.

Global Param Profile

- **Profile Name** - User-configured Global Parameter Profile name.
- **DSCP Number** - The SIP Snooping DSCP number. By default, the packet gets its priority as normal packet (Range = 0 to 63). If the DSCP Number field is set to "0", the value is set as "NA" on switch.
- **SOS Call DSCP No.** - The SIP Snooping SOS Call DSCP No. (Range = 0 to 63)
- **Threshold No. of Calls** - The Number of call records that can be stored in flash (Range = 50 to 500)
- **Clear Stats** - If set to "Yes", when the profile is assigned, existing SIP Statistics are cleared.
- **Reserved Hardware Resource** - Reserved hardware resources required to program ACLs for media entries. Each value is a multiple of the default reserved hardware resources.

- **SIP CPU Rate Limit** - The rate limit of SIP PDUs trapping toward CPU (not applicable to SIP PDUs going towards network port).

Trusted Server Profile

- **Profile Name** - User-configured Trusted Server Profile name.
- **IP Address List** - IP addresses of Trusted Servers included in the profile.

Threshold Profile

- **Profile Name** - User-configured profile name.
- **Jitter** - The Jitter Threshold, in milliseconds (Range = 0 to 300, Defaults = Audio - 50, Video - 100, Other - 100).
- **Packet Loss** - The Packet Loss Threshold, in % (Range = 0 to 99, Defaults = Audio - 10, Video - 20, Other - 20).
- **Round Trip Delay** - The Round Trip Delay Threshold, in milliseconds (Range = 0 to 500, Defaults = Audio - 80, Video - 250, Other - 250).
- **R Factor** - The R-Factor Threshold, in milliseconds (Range = 0 - 100, Defaults = Audio - 70, Video - 80, Other - 80).
- **MOS** - The MOS Value Threshold (Range = 0 - 5, Defaults = Audio - 3.6, Video - 3.0, Other - 3.0). Note that the previous MOS range was 0 - 50. The current range of 0 - 5 represents 1/10th of the previous values.

SOS Profile

- **Profile Name** - User-configured profile name.
- **SOS Call Number List** - The SOS number(s) configured for the profile.


SIP Statistics

- **Total calls processed** - Total calls processed for SIP Snooping.
- **Total audio streams** - Total audio streams.
- **Total video streams** - Total video streams.
- **Total other streams** - Total other streams.
- **Total audio streams that crossed threshold** - Total audio streams that exceeded threshold.
- **Total video streams that crossed threshold** - Total video streams that exceeded threshold.
- **Total other streams that crossed threshold** - Total other streams that exceeded threshold.
- **Number of active calls** - Number of active calls.
- **Number of active audio streams** - Number of active audio streams.
- **Number of active video streams** - Number of active video streams.
- **Number of active other streams** - Number of active other streams.
- **Number of SIP packet received by hardware** - Number of SIP packets received by hardware.
- **Number of SIP packet received by software** - Number of SIP packets received by software.
- **Number of SIP packet received by per method** - Method by which the SIP packet was received (Invite, Ack, Bye, Update and Prack).
- **Number of SIP response packet received** - Number of SIP response packets received.
- **Number of discarded/malformed/unsupported SIP packets** - Number of discarded, malformed or unsupported SIP packets.

- **Number of discarded SIP packets not from/to trusted servers** - Number of discarded SIP packets not from or to trusted servers.
- **Number of dropped SIP packets due the software error** - Number of SIP packets dropped due the software error. (e.g., NI overflow, NI/CMM, CMM overflow).
- **Total Emergency Calls** - Total number of Emergency Calls.

SIP Settings

The SIP Settings Screen is used to enable/disable and configure data retention parameters for SIP data. Configure a field and click on the **Apply** button. The change takes effect immediately. Note that you can click on **Revert** button before applying a change to return a field to its previous value. The definitions below apply to both Detailed and Aggregate Records Data.

- **Exclude Authorization Call** - Use the slider to configure whether or not to include (On) or exclude (Off) Authorization Calls when collecting SIP data.
- **Enable Data Retention Policy** - Use the slider to Enable/Disable data retention. When Data Retention is enabled, information that is older than the number of days specified in the "Days To Retain" field is automatically deleted from the database. In addition, the number of days the data has been retained is listed. If Data Retention is disabled, OmniVista will not remove detailed/aggregated data from database and will accumulate unbounded data. To purge all call record data, click the **Purge All** button, then click **Yes** at the confirmation prompt. To refresh the data display with the most recent information, click on the Refresh icon .
- **Days to Retain** - Set the number of days to retain SIP Data. Default for Detailed Data is 30 Days. Default for Aggregate Data is 365 Days.
- **Number of Records** - The total number of records since the last data purge. Note that when the number of records exceeds 100,000, the color of the "Number of Records" text will change to RED to indicate that a change in database retention policy should be made to reduce the number of records.
- **Days Retained** - The number of days the data has been retained since the last data purge.

23.0 Topology

The Topology application enables you to view the topology of devices in the network, view information about a specific device and perform certain operations on those devices (e.g., edit a device, telnet to a device, reboot a device). You can view devices in a topology map in various ways. For example, you can view all discovered devices in the Physical Network Map (default), and you can create custom maps that enable you to group and display devices in a way that is meaningful for your individual network configuration. You can also highlight specific devices or links, and re-arrange devices in a map and save that new map view. The figure below provides an overview of some of the functions that can be performed in the Topology application. Specific functions available when working with maps are detailed below.

Note: Devices must first be discovered using the Discovery application before they are displayed in Topology.

Topology

The screenshot displays the Alcatel-Lucent Enterprise Topology application. The interface includes a navigation bar with 'Home', 'admin', 'Help', 'Videos', 'About', and 'Logout'. Below the navigation bar, there are tabs for 'Physical Network', 'Current Map', and a search bar. The main area shows a network topology map with devices represented by icons and IP addresses (e.g., 10.255.225.233, 10.255.225.237, 10.255.225.153, 10.255.225.252, 10.255.225.242, 10.255.225.232, 10.255.221.239, 10.255.225.97, 10.255.221.101, 10.255.225.230). A 'Device Status' legend indicates that green represents 'Up'. A 'Device Trap Status' legend shows a green circle for 'Up' and an orange circle for 'Down'. The interface also features a 'Detail Panel' on the right showing device information (Address, MAC Address, Serial Number, Type, Status, Version, Location, System Contact) and an 'Operations' menu (Copy Device to Map, Discovery, Edit Device, Delete Device, Copy as New Device, Ping Device, Poll Device, Poll Link). Annotations provide instructions on how to interact with the map and the application's features.

The Physical Network Map (shown above) is the default map view. It displays all network devices. The devices/maps you can see depend on your Role and permissions as configured in the Users and User Groups application (Security - Users and User Groups). You can also create Custom Maps, or configure Dynamic Logical Maps. These maps are logical maps created from devices in the Physical Network Map.

Note: You can click on the **Go To Table View** button at the top of the screen to go to the Discovery Inventory Screen to see a table view of all discovered devices.

Topology Maps

The Topology application not only provides an overview of the network, it can also be used to perform many functions that you can use to view and configure network devices. These functions are detailed in the following sections:

- Working with Topology Maps
 - Viewing Map Information
 - Devices
 - Links
 - Map Clusters
 - Customizing Maps
 - Customizing Map Colors
 - Changing the Map Layout
 - Selecting Devices
 - Highlighting Devices/Links/Alarms
 - Searching for Devices/Links
 - Creating/Cloning/Editing/Deleting Maps
 - Viewing Different Maps
- Working with Network Devices
 - Viewing Device Information
 - Performing Device Operations
 - Creating a New Device

Working with Topology Maps



















There are a number of options for viewing Topology maps in OmniVista. You can highlight devices or links by different criteria (e.g., device type, device admin state, link admin state, alarm severity), change the map layout, search for specific devices in the map, or create custom logical maps. Note that the maps that you can view and the tasks you can perform in Topology (e.g., creating maps, creating devices) depend on your User Role and User Group (e.g., Administrator, Network Administrator) as defined in the Users and User Groups application).

Viewing Map Information

Topology maps provide a visual representation of the network. The administrative status of the device or link is indicated by color (described below), and detailed information can be displayed by clicking on device or link.

Devices

Each device type is displayed with a unique symbol as shown in the table below. Device status (e.g., Up, Warning, Down) is displayed by the device status circle around the device. Notifications status is displayed in a small circle in the upper-right corner of the circle. This indicates the status of device traps received (e.g., Normal, Minor). Note that this Notifications status is only displayed if traps have been configured on the device. Click on a device to view detailed information about the device (e.g., device type, software version).

Device Icons		
 OS6200	 OS6855	 OS10K
 OS6250	 OS6860/6860E	 Wireless Controller
 OS6350	 OS6900	 Access Point
 OS6400	 OS9000/9000E	 IAP
 OS6450	 OS9700/9700E	 Cisco
 OS6850/6850E	 OS9900	 Third-Party Device

Device Status

Device status is displayed by the device status circle around the device.

- **Green** = Up (Device is up)
- **Orange** = Warning (indicates that traps have been received on the device. The highest level of trap received by the device is displayed (Green, Orange, Red) in the Notifications Status).
- **Red** = Down (Device is down)

Note: The colors above are the default Device Status colors. You can change the colors using the Network Status Screen in the Preferences application (Preferences - User Settings - Colors - Network Status).

Notifications Status

Notifications status displayed in the small circle in the upper right corner of the device, indicating the highest level of trap received by the device:

- **No Circle** = Alarm status is Normal.
- **Green** = Alarm status is Warning.
- **Purple** = Alarm status is Minor.
- **Yellow** = Alarm status is Major.
- **Red** = Alarm status is Critical.

Note: The colors above are the default Notifications colors. You can change the colors using the Alarms Screen in the Preferences application (Preferences - User Settings - Colors - Alarms).

Links

Links between devices are displayed as a single line, whether there is a single link or multiple links. "Green" indicates the link is up. (If there are multiple links, "Green" indicates **all** of the links are up.) "Red" indicates the link is down. (If there are multiple links and even **one** of the links is down, the link will display Red.)

To display detailed link information, move the mouse over the link until the pointer turns into a finger and click. Link information is displayed on the right side of the screen as shown below. If there is a single link, the information is displayed for that link (as shown on the left below). If there are multiple links, a list of the links is displayed (as shown on the right). Click on a link to display details for the link. Click on the "Back" link to return to the list of links.

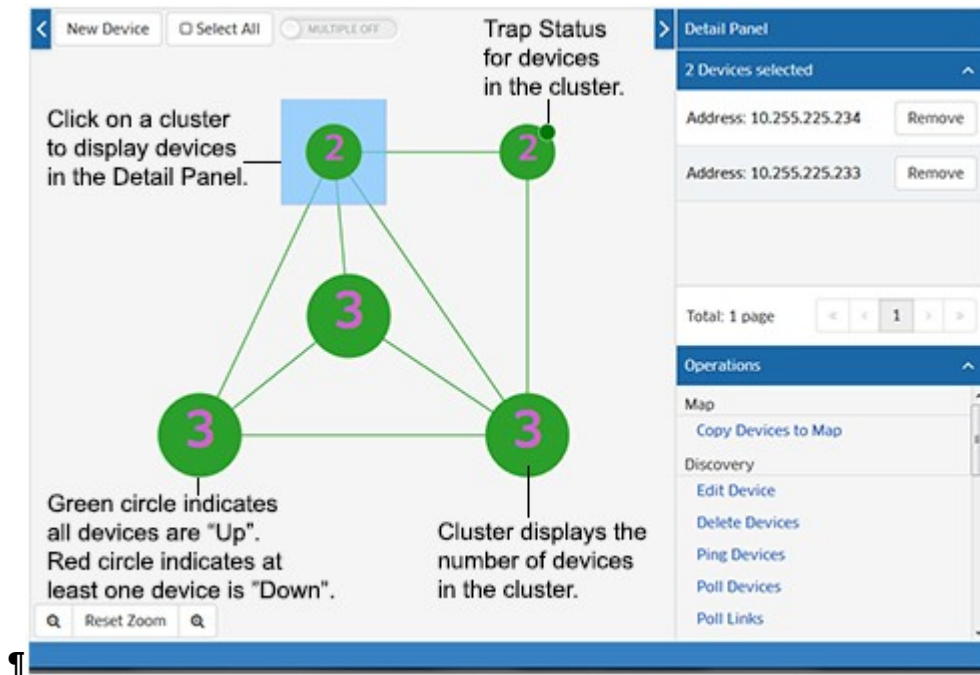
1 Link selected		4 Links selected	
Origin	LLDP	Local IP: 10.255.82.45 - 1/1	Remote IP: 10.255.82.1 - 1/5/9
Local IP	10.255.82.71	Local IP: 10.255.82.45 - 8/5	Remote IP: 10.255.82.1 - 1/3/9
Local Port	1/1/24	Local IP: 10.255.82.45 - 8/6	Remote IP: 10.255.82.1 - 1/4/9
Local LAG Id		Local IP: 10.255.82.45 - 1/2	Remote IP: 10.255.82.1 - 1/6/9
Remote IP	10.255.82.5		
Remote Port	1/3/24		
Remote LAG Id			
Ring Id			
Type	ETHERNET-CRMA/CD		
Status	Up		

The information displayed varies depending on the link type.

- **Origin** - The origin of the link (LLDP, AMAP, Manual).
- **Local IP** - The IP address of the first device in the link.
- **Local Port** - The slot and port that connect the link on the first device, specified above.
- **Local LAG ID** - If this is a link aggregation link, this field displays the Link Aggregation reference number assigned by the first device when the link aggregation group was created.
- **Remote IP** - The IP address of the second device in the link.
- **Remote LAG ID** - If this is a link aggregation link, this field displays the Link Aggregation reference number assigned by the second device when the link aggregation group was created.
- **Type** - The media type of the link.
- **Status** - The status of the link (Up, Down, Unknown).

Map Clusters

When you zoom out for an overview of the network, devices are grouped into clusters, as shown below. Devices are grouped based on their proximity to each other in the map. Each cluster displays the status of devices in the cluster and the number of devices in the cluster, along with a Trap Status indicator at the top of the cluster.



If all devices in a cluster are "Up", the cluster circle displays in Green. If any device in the cluster is "Down", the cluster circle displays in Red. Click on a cluster to display all of the devices in the cluster in the Detail Panel. You can then click on a device to view device details or perform an operation on the device.

The small circle at the top of the cluster displays the status of the highest severity trap generated by any device in the cluster. For example, if the highest level severity trap for any device in the cluster is "Warning", the circle will be Green. If the highest level severity trap for any device in the cluster is "Critical", the circle will be Red. If all traps generated by all devices in a cluster are "Normal", the Trap Status circle will not be displayed. The default trap status colors are shown below.


- **No Circle** = Alarm status is Normal.
- **Green** = Alarm status is Warning.
- **Purple** = Alarm status is Minor.
- **Yellow** = Alarm status is Major.
- **Red** = Alarm status is Critical.

Note: You can change the default colors using the Alarms Screen in the Preferences application (Preferences - User Settings - Colors - Alarms).

Customizing Maps

You can customize the map display (e.g., map background, device/link status colors) and map layout.

Customizing Map Display

You can change the default map preferences (e.g., map colors, link priority) by clicking on the Configuration icon  at the top right corner of the map to bring up the Configuration Window (shown below).

- **Grid Settings** - Show or Hide the Topology Map Grid. When you move devices on a map, they snap to gridlines in the map to enable you to be more precise in your placement of the device in the map. By default, the gridlines are not shown. move the slider to Show, to display the gridlines on the map. Note that showing the gridlines will impact performance. It is recommended that the gridlines be shown only for arranging devices.
- **Favorite Map** - Used to set the default map displayed when you log into OmniVista. Select a map from the drop-down menu to set the default map. By default, the Physical Network Map is displayed.
- **List Map Sort Order** - Used to set the default list order for maps in the Map drop-down menu alphabetically in ascending (A-Z) or descending (Z-A) order. You can always temporarily change the list order by clicking on the sort order in the Map drop-down menu. The order will return to the default setting when you navigate away from the Topology application.
- **Color** - Click on the arrow next to a field to configure a different color, click **Choose**, then click on the **Save** button. The changes take effect immediately.
 - **Device Label Color** - The color of the label under a device (e.g., IP address).
 - **Map Background Color** - The color of the background for all maps.
 - **Device External Link Icon Color** - The color of the device external link icon. If a device in a map has a link to a device that is not in the currently-displayed map, a small arrow (🔗) is displayed in the lower right corner of the device.
- **Link Priority** - Used to set the order in which links are displayed in the Details Panel when a multiple link is selected on a map.


Note: You can change the information e.g., IP address, Device Name, DNS Name) displayed under device on a map using the Device Naming Screen in the Preferences application (Administrator - Preferences - User Settings - Device Naming). This changes the how devices are identified and displayed in all applications in OmniVista.

Changing the Map Layout

You can change the basic layout of a map by clicking on one of the Map Layout options in the lower left corner of the screen, as described below:

- **Circular Layout** - Arranges all of the devices in a circular pattern.
- **Aligned Layout (Default)** - Arranges all of the devices in rows.

- **Circularize Connected Nodes** - Select a device in the map and click on this option to display all nodes connected to the selected device in a circular pattern.

You can also change the layout of a map by selecting a device and dragging it to a new location. You can also select multiple devices and move them. When moving a device(s), the device will snap to a built-in grid to enable you to be more precise in your placement of the device(s) in the map. Click on the Save icon  at the top of the map to save the layout changes. You can also zoom in or out on a map using the Zoom control at the bottom of the map; and you can move an entire map on the screen by clicking anywhere outside the map and dragging the map to a new location on the screen.

Note: If try to leave a map view or leave the Topology application without saving the new layout, you will be prompted to save it. Click **Save** at the confirmation prompt to save the new layout. The map will appear with the new layout the next time it is opened.

Selecting Devices

You can select a single device in a map to perform an operation by clicking on the device. If you want to select multiple devices for an operation, enable the **Multiple On** slider at the top of the map and click and drag to select a group of devices. The devices will appear in a list in the Detail Panel on the right side of the screen (click on the **Remove** button to remove a device from the list). You can then select an operation (e.g., Copy Devices to Map, Poll for Traps, Reboot) to perform on the devices. Note that the operations available depend on the device type(s) selected and whether or not you select a single device or multiple devices. You can also select all devices in a map by clicking on the **Select All** checkbox at the top of the map.

Highlighting Devices/Links/Alarms

You can highlight a devices or links in a map by selecting one or more of Highlight criteria on the left side of the screen (Highlight Device, Highlight Link, Highlight Alarm). When you select one or more of the criteria (e.g., Status Up, Stack, Critical Alarm), only those devices/links matching the all of the selected criteria are displayed in the map. To return to the original map view, click on the selection criteria again to de-select it. You can highlight devices/links based on the following criteria.

- **Highlight Device**
 - Status
 - Up
 - Down
 - Warning
 - Type
 - Stack
 - Virtual Chassis
 - OAW
 - Configuration
 - Need Certify
 - Unsaved
 - Synchronization
 - Need Synchronize
- **Highlight Link**
 - Status
 - Up
 - Down
 - Type
 - Aggregate

- Manual Link
- Discovered Link
- **Highlight Alarm**
 - Critical
 - Major
 - Minor
 - Warning
 - Normal

Searching for Devices/Links

You can search for a device(s) or link(s) in the currently displayed map, or search in all maps. At the top of the map, select "Current Map" or "All Maps" from the drop-down menu. Click on the "Search" area to see a list of search criteria that you can search on (e.g., IP address, MAC address, Link Title, Local IP). Enter the search information. Any device(s)/link(s) matching the criteria will be displayed in a list below the Search area. Click on an item to highlight the device/link in the map. Click on the x in the "Search" area to delete the search criteria and begin a new search.

Creating/Cloning/Editing/Deleting Maps

The Physical Network Map view (Default) displays all network devices. You can also create Custom Maps, or configure Dynamic Logical Maps. These maps are logical maps created from devices in the Physical Network Map. You can create multiple maps (Custom and Dynamic), and a device can be included in multiple maps.

Note: Admin and Netadmin Users can create/edit/delete maps.

Creating Maps

Click on the **New Map** button at the top of the screen. The Create Custom Map window will appear. Complete the fields as described below to create either a Custom Map or a Dynamic Map. When you are finished, click on the **Create** button to create the map.

- **Map Name** - Enter a unique name for the map. You cannot duplicate the name of an existing map.
- **Background** - If you want to change the default background for the map, select an image option, then select the background from the drop-down menu:
 - **Upload New** - Click on the **Browse** button to locate a new background image, then select the image from the drop-down menu.
 - **Existing Images** - Select an existing image from the drop-down menu.
- **Dynamic Map (to create a Dynamic Map)** - You can use a filter to create a map that will dynamically add/remove devices from a map. Set the **Dynamic Map** slider to "On". Click on the **Filter** button to bring up the Filter Selection Window, and select a filter to apply to the map. If necessary, click on the Create icon **+** to create a new filter. (You can create a filter for IP address, Device Type, AOS Version, or Device Location.) Devices will now be added to/removed from the map dynamically based on the applied filter.
- **Devices (to create a Custom Map)** - You can also create a Custom Map to include specific devices. Use the switch picker to select the devices you want to include the Custom Map.

Note: You can also copy devices from one map to another using the "Copy Device to Map" operation. See Performing Device Operations for more information.


Note: If there are performance issues in rendering larger maps, you can reduce the

number of devices in the map to improve performance.


Cloning Maps

You can also "clone" an existing map to quickly create a new Custom Map. Select a map that you want to clone from the drop-down menu at the top of the screen. Click on the **Clone Map** button. The Clone Map window will appear with all of the devices in the map you are cloning pre-selected. Enter a **Map Name** and complete the fields as described above to create a new Custom Map.


Editing Maps

Click on the Edit icon  in the upper-right corner of the screen and edit any fields as described above. You cannot edit the map name. You cannot change a Custom Map to a Dynamic Map or a Dynamic Map to a Custom Map; or edit the filter for a Dynamic Map.

Deleting Maps

Go to the map you want to delete by selecting it from the Map drop-down menu at the top-left corner of the screen. When the map appears, click on the Delete icon  in the upper-right corner of the screen, then click **OK** at the confirmation prompt. Note that you cannot delete the Physical Network Map.

Viewing Different Maps

The Map drop-down menu in the upper-left corner of the screen displays the name of the current map. To change to a different map, select the map from the drop-down menu. You can also sort the list of maps in ascending or descending order or search for a specific map in the list. Click on the Map drop-down and then click on the up or down arrows to sort, or enter text in the Search field to search. If you change to sort order, the sort will remain persistent for as long as you are in the Topology application. Once you navigate to another application, the order will revert to the configured default order. You can always re-sort the maps when you return to Topology by clicking on the Configuration icon  at the top right corner of the map to bring up the Configuration Window.

Working with Network Devices

You can view detailed information about a device or perform certain operations on a device by clicking on the device in the map. Operations can be performed on a single device or multiple devices. See *Selecting Devices* for instructions on selecting single or multiple devices. After selecting a device(s), you can perform one of the operations displayed in the Operations area. Note that the operations available depend on the device type(s) selected and whether or not you select a single device or multiple devices.

Note: The operations you can perform in Topology (e.g., performing backups, creating inventory reports) depend on your User Role and User Group configured in the Users and User Groups application.

Viewing Device Information

Click on a device to display detailed device information. The information displayed may vary depending on the device. You can only view information for a single device. Device information is not displayed if you select multiple devices.

- **Address** - The IP address of the device.
- **MAC Address** - The MAC address of the device.
- **Serial Number** - The serial number of the device.
- **Type** - The device type/model (e.g., OS6900-X20)

- **Status** - The administrative status of the device:
 - **Up** - Device is up and responding to polls.
 - **Warning** - Device has sent at least one warning or critical trap and is thus in the warning state.
 - **Down** - Device is down and not responding to polls.
- **Version** - The software version installed on the device (e.g., 7.x.x.317.R01). OmniVista may not be able to determine the software version on some third-party devices. In these cases, the field will be blank.
- **Location** - The physical location of the device.
- **System Contact** - Contact information for the person responsible for the device.
- **Name** - The device name.
- **DNS Name** - The device DNS name, if applicable.
- **Last Upgrade Status** - The status of the last firmware upgrade on the switch:
 - **Successful** - Successful BMF and Image upgrade performed.
 - **Successful (BMF)** - Successful BMF upgrade performed.
 - **Successful (Image)** - Successful Image upgrade is performed.
 - **Failed (BMF, Image)** - BMF and Image upgrade failed.
 - **Failed (BMF)** - BMF upgrade failed.
 - **Failed (Image)** - Image upgrade failed.
- **Backup Date** - The date that the device's configuration and/or image files were last backed-up to the OmniVista Server.
- **Backup Version** - The firmware version of the configuration and/or image files that were last backed-up to the OmniVista Server.
- **Last Known Up At** - The date and time when the last poll was initiated on the device.
- **Description** - A description of the device, usually the vendor name and model.
- **Traps** - The status of trap configuration for the device
 - **On** - Traps are enabled.
 - **Off** - Traps are disabled.
 - **Not Configurable** - Traps for this device are not configurable from OmniVista. (Note that traps may have been configured for such devices outside of OmniVista.) **Unknown** - OmniVista does not know the status of trap configuration on this device. OmniVista will read the switch's trap configuration when traps are configured for the switch via the Configure Traps Wizard.
- **Number of Licenses** - The total number of Core (AOS) or Third-Party licenses being used. For example, a stack of 4 switches would require 4 licenses, a VC of 6 would require 6 licenses. If a stack splits, the number of licenses reserved for the device before the split is maintained even though modules have been reduced to less than 5. This way, the license counts are reserved for the stack to recover.
- **License Type** - The type of license used by the device (e.g., AOS, ThirdParty).
- **Running From** - For AOS devices, this field indicates whether the switch is running from the **certified** directory or the **working** directory. This field is blank for all other devices. For AOS devices, the directory structure that stores the switch's image and configuration files in flash memory is divided into two parts:
 - **Certified Directory** - Contains files that have been certified by an authorized user as the default configuration files for the switch. When the switch reboots, it will automatically load its configuration files from the certified directory if the switch detects a difference between the certified directory and the working directory.
 - **Working Directory** - Contains files that may or may not have been altered from those in the certified directory. The working directory is a holding place for new files to be tested before committing the files to the certified directory. You can save configuration changes to the working directory. You cannot save configuration changes directly to the certified directory.

Note that the files in the certified directory and in the working directory may be different from the running configuration of the switch, which is contained in RAM. The running configuration is the

current operating parameters of the switch, which are originally loaded from the certified or working directory but may have been modified through CLI commands, WebView commands, or OmniVista. Modifications made to the running configuration must be saved to the working directory (or lost). The working directory can then be copied to the certified directory if and when desired.

Note: OmniVista supports the Multiple Working Directories Feature available on OS10K and OS6900 Switches (AOS Release 7.2.1.R01 and later). This feature allows the user to create multiple Working Directories on the switch that can be used to save specific switch configurations. The user can create any name for these "Working" Directories (e.g., "Marketing Switch 05-10-17"). If the switch is running from one of these user-created directories, the directory name is displayed in this field.

- **Changes** - For AOS devices, this field indicates the state of changes made to the switch's configuration. This field is blank for all other devices. Information in the Changes field will be accurate as long as OmniVista has polled the switch since the last change was made (through any interface). Note that it is possible a switch could be in a state where it is both Unsaved and Uncertified. In this situation "Unsaved" displays in the Changes field. This field can display the following values:
 - **Certified** - Changes have been saved to the working directory, and the working directory has been copied to the certified directory. The working directory and the certified directory are thus identical.
 - **Uncertified** - Changes have been saved to the working directory, but the working directory has not been copied to the certified directory. The working directory and the certified directory are thus different.
 - **Unsaved** - Changes have been made to the running configuration of the switch, but they have not been saved to the working directory.
 - **Blank** - When this field is blank for an AOS device, it means that OmniVista knows of no unsaved configuration changes and assumes that the working and certified directories in flash memory are identical.
- **Synchronized Status** - Indicates whether the primary CMM module's working directory is identical to the working directory on the other CMM module (if present):
 - **Synchronized** - The primary CMM module's working directory is identical to the working directory on the other CMM module.
 - **Need Synchronize** - The primary CMM module's working directory is not identical to the working directory on the other CMM module.
 - **Not Applicable** - Only one CMM module is installed.
 - **Unknown** - The synchronization state is unknown.

Performing Device Operations

You can perform the following operations on any device in a map. To execute one of the operations, click on the device and select an operation from the Operations area in the lower-right corner of the screen. Note that the operations available are different for single device selection and multiple device selection. Available operations will also vary if you select different types of devices in multiple selection mode.

Map

- **Overlay View** - Provides detailed information for wireless, virtual chassis, and stack devices. Topology maps display a single icon for Virtual Chassis, Stacks, and Wireless devices. For example, topology maps display only the Master Chassis in a virtual chassis. You can select the device in the map and click on "Overlay View" to display the other the virtual chassis or stack.

When the overlay view is displayed, you can click on any device in the overlay view for detailed information on the device. The following information is displayed:

- **Virtual Chassis Overlay** - Displays the devices and links between them. Click a node/link to view chassis/VFL link detail information.
 - Chassis Node - Oper Chas ID, Role, Chas Status, Oper Status, Cfg Chas ID, Oper Priority, Group, MAC Address, Model.
 - VFL Link - Local Port, Remote Port, Link Status.
- **Stack Overlay** - Displays each slot (chassis), and the links between them. Click on a node/link to view slot/link detail information.
 - Slot - Slot NI, Role, Slot Status, Oper Status, Saved NI Number, MAC Address, Model.
 - Link - Local Slot NI, Local Port, Remote Slot NI, Remote Port, Link Status.
- **Wireless Device Overlay** - Displays logical links between Controller and Access Points (APs). The following information is displayed:
 - Controller Nodes - Controller IP, MAC Address, Controller Status, Master IP (IP address of master AP), Model, Version, Name, Location, DNS Name.
 - AP Nodes - AP MAC Address, AP IP Address, AP Name, AP Model, AP Location, AP Status.
- **Copy Device to Map** - Copies the selected device(s) from the current map to a different map. Select a device(s), click on the "Copy Device to Map" operation link, then select the new map location from the Map drop-down at the top-left of the screen. When the new map appears, click on the **Add devices to this map** button. The device(s) will appear in the new map.

Discovery

- **Edit Device** - Opens up the Edit Discovery Manager Screen in the Discovery application for the selected device.
- **Delete Device** - Deletes the selected device from the current map.
- **Copy as New Device** - Takes you to the Discovery application's "Add New Device" Screen, where you can add a new device based on the configuration (e.g., CLI/FTP Password, SNMP configuration) of selected device.
- **Ping Device** - Causes an immediate ping of the selected device(s), and launches the Discovery Screen in the Discovery Application to display the results.
- **Poll Device** - Causes an immediate poll of the selected device(s), and launches the Discovery Screen in the Discovery Application to display the results.
- **Poll Link** - Causes an immediate poll of any links on the device(s) switches, and launches the Discovery Screen in the Discovery Application to display the results.
- **Reboot** - Reboots the selected device(s) You have the option of rebooting from the Working, Certified, or Other Directory and setting a time for the reboot. Click on the Reboot operation link and use the **Reboot From** drop-down to select the directory you want to reboot from. In the **Reboot Delay** drop-down select when you want to reboot to occur (now, a specific number of minutes from now, or at a specific date and time). Note that when you reboot multiple devices, there is a minimum delay of 30 seconds before the devices reboot (even if you select the Reboot now option). If you select a large number of devices, the delay is equal to roundoff of $(30 + (\text{deviceCount}/4))$, in seconds (e.g., if you select 1,000 devices, the delay is 280 seconds, or 4 minutes). The delay allows time to push the "Reboot" command to all devices.
- **Copy Running/Working to Certified** - Copies the contents of the working/running directory in the primary CMM to the certified directory in the primary CMM. Note that the Copy Working to Certified command also automatically synchronizes the switch's CMMs after the copy operation is completed.
- **Copy Certified to Working/Running** - Copies the contents of the certified directory in the primary CMM to the working/running directory in the primary CMM.
- **Save to Running** - Saves the primary CMM's current running configuration to the current

running directory of the switch. OmniVista supports the Multiple Working Directories Feature on certain devices (e.g., OS10K, OS6900). This feature allows the user to create multiple "working" directories on the switch that can be used to save specific switch configurations. When the Save to Running Command is executed, the device(s) save the CMM's current running directory to the current user- defined "working" directory (Running Directory). Note that if you select a group of devices and some do not support multiple working directories, the devices will save the CMM's current running directory to the device's current "working" directory, whether it is a user-defined directory or the Working Directory.

- **Webpage** - Opens up a Web session with the selected device. The web session application varies depending on the device. For example, AOS devices will open a WebView session.
- **Configure Health Thresholds** - Opens the Configure Device Health Thresholds Screen in the Discovery application for the selected device. Thresholds are used to set limits for health traps. If a device has been configured to send health traps, a trap will be sent whenever a monitored item's current utilization exceeds the configured health threshold. Configure the CPU, Memory, or Temperature Threshold for the selected device(s) and click on the **Apply** button. Note that you cannot configure the Temperature Threshold. The Temperature Threshold is hard coded on devices.

CLI Scripting

- **SSH** - Opens up a Telnet (SSH) session with the selected device in the CLI Scripting application.

Locator

- **Locate End Stations** - Launches the Locator application and searches for all end stations that are attached to the selected switch. All end stations found are displayed in the Locator application's Browse Screen.

Notifications

- **View Traps** - Launches the Notifications Home Screen in the Notifications Application to display traps for the selected device.
- **Poll for Traps** - Causes an immediate poll of the selected devices for traps. Traps are reported in the Notifications application.
- **Configure Traps** - Launches the Trap Configuration Wizard in the Notifications application to enable you to configure traps for the selected devices.

Resource Manager

- **Backup Device** - Launches the Backup Wizard in the Resource Manager Application, which enables you to perform a configuration backup of the selected devices.
- **Device Inventory** - Launches the Inventory Screen in the Resource Manager Application for the selected switches, which enables you to create and Inventory Report for the selected devices.

Creating a New Device

You can create a new device and add it to the map by clicking on the **New Device** button at the top of the screen. This will take you to the Discovery Application's "Add New Device" Screen, where you can add a new device to the map.

24.0 Unified Access

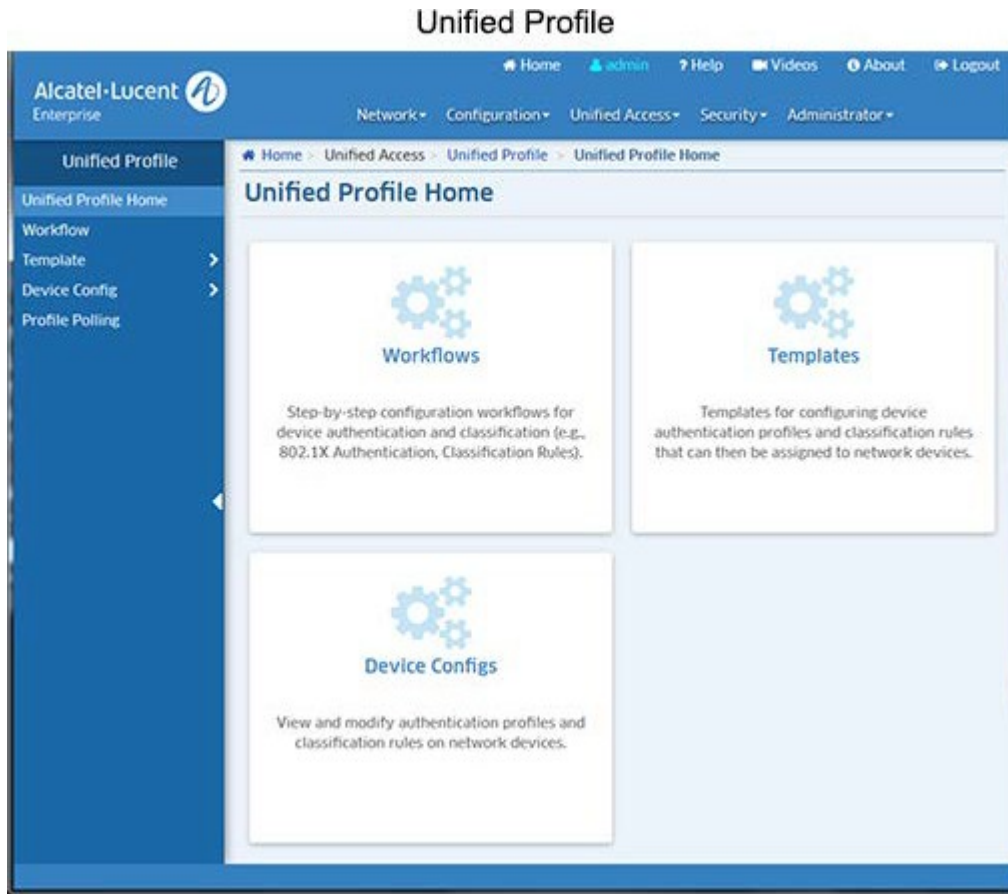
The Unified Access Application provides unified wireline/wireless device configuration of security functions for Edge Ports and Access Points (APs) on OmniAccess Wireless devices. In addition to device authentication and classification, you can create Access Role Profiles (similar to User Network Profiles) to configure network access controls for one or more user devices. Unified Access contains applications that work together with the Authentication Servers application to seamlessly authenticate and configure QoS for both device types. It is configured using the following applications:

- Unified Profile - Provides unified security functions for Edge Ports for supported wireline and wireless devices.
- Unified Policy - Used to configure Unified Policies and Policy Lists, which enable the user to configure QoS policies for both wireline and wireless devices.
- Multimedia Services (mDNS) - Used to configure the Multicast Domain Name System (mDNS) protocol. mDNS is used by "Zero Configuration Networking" solutions such as Apple's Bonjour, Avahi LGPL, and Linux NSS-mDNS. mDNS is a resolution service that is used to discover services on a LAN.
- Premium Services (BYOD) - Used to configure the Bring Your Own Device (BYOD) feature. BYOD leverages Access Guardian features along with ClearPass Policy Manager (CPPM) to allow a wired or wireless guest, device, or authenticated user to connect to the network through an OmniSwitch edge device using the CPPM Server for unified authentication.

Unified Profile

The Unified Access Unified Profile application provides unified security functions for Edge Ports and AOS WLAN Devices. In addition to device authentication and classification, you can create Access Role Profiles (similar to User Network Profiles) to configure network access controls for one or more user devices. This is achieved using both Layer 2 and Layer 3 Authentication and Classification. Layer 2 Authentication and Classification provides the initial user authentication and Access Role Profile assignment. Layer 3 Authentication and Classification can dynamically change the QoS Policy List/Role for a user already authenticated and classified into the network. Based on the Access Role Profile (UNP) into which the user is initially classified, the user may undergo Quarantine Manager and Remediation (QMR), RADIUS based MAC Check Blacklisting, and Location or Time based validations that can restrict a user's network access or assign different Policy Lists/Roles to the user.

Unified Profile is configured using the links on the left side of the screen. An overview of each configuration function is provided below.



Authentication and Classification

Unified Profile provides network access and Quality of Service on a per user basis. This is achieved using both Layer 2 and Layer 3 Authentication and Classification.

Layer 2 Authentication and Classification

Unlike regular Layer 2 bridging, all users are learned through software. The user is learned "forwarding" or "filtering" based on Layer 2 authentication mechanisms configured on the port. The authentication mechanisms supported are 802.1x and MAC based authentication. Access Classification rules can also be used to learn a user in the forwarding state if no authentication mechanism is configured. Apart from determining the forwarding or filtering state, this stage also determines the UNP and VLAN to be assigned to the user. The UNP and VLAN assigned to the user do not change. The UNP provides an initial QoS Policy list/role to be assigned to the user.

A user is first authenticated using 802.1x or MAC based authentication (MAC authentication is used only if 802.1x authentication is disabled or the user is not a supplicant.) If the user passes authentication, and the RADIUS server returns a valid UNP name, the user is mapped to that Access Role Profile (UNP name) and VLAN. The RADIUS server may also return an explicit policy list name, which overrides the policy list associated with the UNP.

If authentication is not enabled or fails, or the Authentication Server does not return a valid UNP, and classification is enabled, the user is classified based on one of the following Access Classification Rules - Port, Group ID, MAC, LLDP, Authentication Type, IP Address - and assigned a Default UNP.

Configuring Unified Profile

Unified Profile can be configured using Workflow windows and/or Templates. Unified Profile is configured using the links on the left side of the screen.

- **Workflow** - OmniVista provides guided workflows for configuring Unified Profile. These workflows can be used for an initial Unified Profile setup, which can then be fine tuned using the Unified Profile Templates and Device Config editing. You can use the workflows to:
 - Classify network traffic based on Access Classification
 - Rules. Use RADIUS Servers to authenticate users using 802.1X.
 - User RADIUS Servers to authenticate users using 802.1X or MAC Address.
 - User RADIUS Servers to authenticate users using MAC Address or Captive Portal
 - Use ClearPass to authenticate users using 802.1X, MAC Address or Web-based credentials.
- **Template** - Unified Profile Templates are used to configure Unified Profiles (e.g., Access Auth Profiles, Wireless Profiles) and apply them to multiple network devices.
- **Device Config** - Device Config screens enable you to view, edit, and delete Unified Profiles on specific network devices.
- **Profile Polling** - Used to set the polling interval at which device configuration information is updated, and perform immediate polls to update Unified Profile information.

Workflow

Unified Profile Workflows are guided workflows you can use to configure Unified Profile. These workflows can be used for setting up initial Unified Profiles on devices for specific use cases, which can then be fine tuned using the Unified Profile Templates and Device Config editing. You can use the workflows to:

- **Traffic Based on Classification Rules** - Classify network traffic based on Access Classification Rules.
- **802.1X Authentication** - Use RADIUS Servers to authenticate users using 802.1X.
- **802.1X and MAC Authentication** - User RADIUS Servers to authenticate users using 802.1X or MAC Address.
- **MAC Authentication and Captive Portal** - User RADIUS Servers to authenticate users using MAC Address or Captive Portal
- **802.1x MAC Authentication and Captive Portal with ClearPass** - Use ClearPass to authenticate users using 802.1X, MAC Address or Web-based credentials. Note that when you apply this workflow to a device, the ClearPass IP address will be automatically configured on the device based on the "unp redirect-server" configured on the device. If you want to change the ClearPass Server, go to the ClearPass Screen (Unified Access - Premium Services - BYOD - ClearPass) to create a new ClearPass Server and apply it to the device.

Unified Profile Templates

Unified Profile Templates are used to configure Unified Access Profiles, which provide unified security functions for Edge Ports and AOS WLAN Devices. In addition to device authentication and classification, you can create Access Role Profiles (similar to User Network Profiles) to configure network access controls for one or more user devices. This is achieved using both Layer 2 and Layer 3 Authentication and Classification. Layer 2 Authentication and Classification provides the initial user authentication and Access Role Profile assignment. Layer 3 Authentication and Classification can dynamically change the QoS Policy List/Role for a user already authenticated and classified into the network. Based on the Access Role Profile (UNP) into which the user is initially classified, the user may undergo Quarantine Manager and Remediation (QMR), RADIUS based MAC Check Blacklisting, and Location or Time based validations that can restrict a user's network access or assign different Policy Lists/Roles to the user.

The first step in configuring Unified Profile is to configure an Access Auth Profile and assign it to ports/linkaggs on the network. You then configure Access Role Profiles and AAA Server Profiles to which a user is assigned based on the Access Auth Profile. The following links are used to access Unified Profile Templates:

- Access Auth Profile - An Access Auth Profile contains all of the UNP properties to be enabled on an Edge Port. The template can be applied to a port or linkagg to enable UNP Edge Port status and set the parameters for the authentication process for the port. The Access Auth Profile configures 802.1x and MAC authentication, Access Classification, the AAA Server Profile to be used for authentication specifying the default Access Role Profile (UNP), etc.
- Wireless Profiles - Used to configure authentication parameters for wireless devices and include those parameters in Unified Access Authentication Profiles.
 - SSID Profile - An SSID Profile can be created and included in an Access Authentication Profile that can be assigned to wireless devices on the network to enable the authentication process for that SSID.
 - AP Group - Used to configure an Access Point (AP) Group on a Controller and then associate the group with APs attached to that controller.
- AAA Server Profile - In addition to the global AAA configuration included in an Access Auth Profile, you can also create a AAA Server Profile that can be applied on a per Port/Linkagg basis. This enables you to configure different RADIUS Servers for different users on different ports and apply different RADIUS client attributes to them. AAA Server Profiles are only supported on 8.x and Wireless devices. 6.x and 7.x devices use Global AAA Configuration.
- Access Role Profile - Contains the various UNP properties, including the QoS Policy List attached to the UNP and Captive Portal Authentication for users assigned to into this UNP.
- Access Classification - Used to configure Access Role Profile Classification Rules.
- Customer Domain - Used to configure Customer Domains. Customer Domains provide an additional method for segregating device traffic. A Customer Domain is identified by a numerical ID, which can be assigned to UNP ports and Access Classification Rules.
- SPB Profile - Used to create an SBP Profile. An SPB Profile contains SBP parameters that can be mapped to an Access Role Profile.
- Far End IP - Used to create Far End IP Lists. A Far End IP List is assigned to an Access Role Profile through the mapping of VXLAN service parameters to the profile. This allows multiple far-end nodes to be associated with the service created for the VXLAN Network ID (VNID) specified in a VXLAN Profile.
- Static Service - Used to configure a Static Service Profile. This can be used to configure the mapping of an existing SPB or VXLAN service ID to an Access Role Profile.
- VXLAN Profile - Used to configure a VXLAN Profile that can be mapped to an Access Role Profile.
- Global Configuration - Used to create global AAA Profiles. Global AAA Profiles are applied in the absence of a template's AAA Profile for 8.x devices.

Access Auth Profile

The Unified Profile Access Auth Profile Screen displays all configured Access Auth Profiles and is used to create, edit, and delete Access Authentication Profiles. An Access Auth Profile enables you to assign a pre-defined UNP port configuration to a port or linkagg, or specify them individually on each port to enable UNP port status and set the parameters for the authentication process for the port. For IAP devices, an Access Auth Profile can be assigned to a WLAN identified by the SSID Profile. For wireless controller devices, an Access Auth Profile can be assigned to Virtual AP Profile, which is used to configure WLAN. The Access Auth Profile configures 802.1X and MAC authentication for both wired and wireless devices, Access Classification and the default AAA Server and/or UNP Profile to be used once a user is authenticated. The basic configuration for each configured Access Auth Profile is displayed. You can also click on a profile for

a configuration view.

Creating an Access Auth Profile

Click on the Create icon **+**. Enter a **Profile Name** and configure the profile as described below, then click on the **Create** button. When you are finished, select the checkbox next to the profile and click on the **Apply to Devices** button to assign the profile to switches/ports or wireless devices/virtual APs on the network.

Default Settings

This section is used to configure basic settings for the profile.

- **AAA Server Profile** - The AAA Server Profile used to authenticate users on the port. Select a profile from the drop-down list or click on the "Add New" link to go to the AAA Server Profile Screen and create a new profile.
- **Port Bounce** - Enables/Disables Port Bounce. Always Enabled on wireless devices and AOS 6x switches. This feature is required to handle scenarios where a client is switched from one VLAN to other after COA. If port bounce is enabled, the port will be administratively shut down. This is to trigger DHCP renewal and re-authentication, if necessary.
- **MAC Auth** - Enables/Disables MAC Authentication for the port. Wireless devices do not contain this attribute in their configuration table. MAC Pass Alt attribute in the next section No Auth/Failure/Alternate is used for MAC Authentication on wireless devices.
- **802.1X Auth** - Enables/Disables 802.1X Authentication. Wireless devices do not contain this attribute in their configuration table. 802.1X Pass Alt attribute in the next section No Auth/Failure/Alternate is used for 802.1X Authentication on wireless devices.
- **Dynamic Service** - Select a dynamic mapping method for the profile, if applicable (SPB, VXLAN).
- **Customer Domain ID** - Select a Customer Domain ID for the profile, if applicable. If necessary, click on the "Add New" link to go to the Customer Domain Screen and create a Customer Domain.

No Auth/Failure/Alternate

This section is used to configure the actions taken if a device assigned to the profile fails authentication.

- **Trust Tag** - Enables/Disables whether or not to trust the VLAN ID of a tagged packet to determine how the packet is classified. Enabling the trust VLAN ID tag option provides an implicit method of VLAN tag classification that will accept tagged traffic without the need to create specific classification rules for those profiles
- **Access Classification** - Enables/Disables device classification. Always Enabled on wireless devices (Default = Disabled).
- **Default Access Role Profile** - The Default Access Role Profile that users are assigned to if authentication or classification methods fail to match traffic with any role. This is the last-resort role. Select a profile from the drop-down list or click on the "Add New" link to go to the Access Role Profile Screen and create a new profile. Note that for IAP devices the default Access Role Profile name must match the SSID Profile name in order for it to take effect.
- **802.1X Pass Alt** - The user shall be assigned a Pass-Alternate UNP in case the 802.1X authentication does not result in a valid UNP for the pass branch. Select a profile from the drop-down list or click on the "Add New" link to go to the Access Role Profile Screen and create a new profile.
- **Bypass Status** - Enables/Disables 802.1X bypass. When 802.1X bypass is enabled, the user's 802.1X authentication method is skipped. The user enters directly mac-authentication or Access Classification based on the configuration on the UNP ports/Linkaggs. On wireless devices, this attribute corresponds to another attribute named I2-auth-fail-through, and this attribute must be combined with the MAC Allow EAP attribute to make I2-auth-fail-through attribute work (Default =

Disabled) .

- Bypass Status with ENABLED status combined with None MAC Allow EAP will disable 802.1X authentication, and I2-auth-fail-through is not ENABLED
- Bypass Status with ENABLED status combined with Fail MAC Allow EAP will enable I2-auth-fail-through.
- Other configurations of Bypass Status and MAC Allow EAP cause I2-auth-fail-through to be ignored on wireless devices.
- **Failure Policy** - The authentication method used if 802.1X authentication fails.
- **MAC Pass Alt** - The Access Role Profile the user is assigned to after passing authentication.
- **MAC Allow EAP** - Enables/Disables Extensible Authentication Protocol (EAP).

Advanced Settings

This section is used to configure advanced 802.1X authentication settings for the profile.

- **802.1X Tx Period Status** - Enables/Disables 802.1X Authentication Tx Period (Default = Disabled).
- **802.1X Tx Period** - Access Auth Profile 802.1X Tx period, in seconds.
- **802.1X Supp Timeout Status** - Enables/Disables 802.1X Supp Timeout (Default = Disabled).
- **802.1X Supp Timeout** - 802.1X Authentication Supp Timeout, in seconds.
- **802.1X Request Status** - Enables/Disables 802.1X Authentication Max Request (Default = Disabled).
- **802.1X Request** - 802.1X Authentication Max Request number.
- **Port Controlled Directions** - Configures whether network access control is applied to both incoming and outgoing traffic, or only applied to incoming traffic (In/Both, Default = Both).

Wireless Settings

This section is used to configure a Virtual AP Profile (i.e., "wireless device" profile) and associate it with the Access Auth Profile.

- **Virtual AP Name** - User-configured name for the Virtual AP Profile.
- **SSID Profile** - The SSID Profile you want to associate with the Virtual AP Profile. Select a profile from the drop-down list or click on the "Add New" link to go to the SSID Profile Screen and create a new profile.
- **User Derivation Rules** - Select a User Derivation Rules from the drop-down list to specify a user attribute profile from which the user role or VLAN is derived. The user role can be derived from user attributes upon the client's association with an AP (this is known as a user-derived role). You can configure rules that assign a user role to clients that match a certain set of criteria. For example, you can configure a rule to assign the role VoIP-Phone to any client that has a MAC address that starts with bytes xx:yy:zz. User-derivation rules are executed before client authentication. Note that only wireless classification rules are listed in the drop-down menu.
- **Virtual AP Enable** - Enables/Disables the Wireless Authentication Profile.
- **Forward Mode** - Controls whether data is tunneled to the controller using generic routing encapsulation (GRE), bridged into the local Ethernet LAN (for remote APs), or using a combination of both depending on the destination (e.g., corporate traffic goes to the controller, and Internet access remains local). All forwarding modes support band steering, TSPEC/TCLAS enforcement, 802.11k and station blacklisting.
 - **Tunnel** - The AP handles all 802.11 association requests and responses, but sends all 802.11 data packets, action frames, and EAPOL frames over a GRE tunnel to the controller for processing. The controller removes or adds the GRE headers, decrypts or encrypts 802.11 frames, and applies firewall rules to the user traffic as usual. Both remote and campus APs can be configured in tunnel mode.
 - **Bridge** - 802.11 frames are bridged into the local Ethernet LAN. When a remote AP or

campus AP is in bridge mode, the AP (and not the controller) handles all 802.11 association requests and responses, encryption/decryption processes, and firewall enforcement. The 802.11e and 802.11k action frames are also processed by the AP, which then sends out responses as needed.

- **Split Tunnel** - 802.11 frames are either tunneled or bridged, depending on the destination (e.g., corporate traffic goes to the controller, and Internet access remains local).
- **Decrypt Tunnel** - Both remote and campus APs can be configured in decrypt-tunnel mode. When an AP uses decrypt-tunnel forwarding mode, that AP decrypts and decapsulates all 802.11 frames from a client and sends the 802.3 frames through the GRE tunnel to the controller, which then applies firewall policies to the user traffic.
- **Allowed Band** - The band(s) on which to use the Virtual AP:
 - **a** - 802.11a band only (5 GHz)
 - **g** - 802.11b/g band only (2.4 GHz)
 - **all** - Both 802.11a and 802.11b/g bands (5 GHz and 2.4 GHz). (Default).
- **Band Steering** - Enables/Disables Band Steering. Band Steering encourages dual-band capable clients to stay on the 5GHz band on dual-band APs. This frees up resources on the 2.4GHz band for single band clients like VoIP phones. The feature supports both campus APs and remote APs that have a virtual AP profile set to tunnel, split-tunnel or bridge forwarding mode. Note, however, that if a campus or remote APs have virtual AP profiles configured in bridge or split-tunnel forwarding mode but no virtual APs in tunnel mode, those APs will gather information about 5G-capable clients independently and will not exchange this information with other APs that only have bridge or split-tunnel virtual APs configured.
- **Steering Mode** - Band steering supports the following three band steering modes.
 - **Force-5GHz** - The AP will try to force 5Ghz-capable APs to use that radio band.
 - **Prefer-5GHz** -The AP will try to steer the client to 5G band (if the client is 5G capable) but will let the client connect on the 2.4G band if the client persists in 2.4G association attempts.
 - (Default) **Band Balancing** - The AP tries to balance the clients across the two radios in order to best utilize the available 2.4G bandwidth. This feature takes into account the fact that the 5Ghz band has more channels than the 2.4 GHz band, and that the 5Ghz channels operate in 40MHz while the 2.5Ghz band operates in 20MHz.
- **Dynamic Multicast Optimization** - Enables/Disables Dynamic Multicast Optimization.
- **Dynamic Multicast Optimization Threshold** - The maximum number of high-throughput stations in a multicast group beyond which dynamic multicast optimization stops. (Range = 2 - 255, Default = 5)
- **Drop All Broadcast or Multicast Traffic** - If "Enabled", broadcast and multicast traffic is dropped. Do not enable this option for Virtual APs configured in bridge forwarding mode. This configuration parameter is only intended for use for Virtual APs in tunnel mode. In tunnel mode, all packets travel to the controller, so the controller is able to drop all broadcast traffic. When a Virtual AP is configured to use bridge forwarding mode, most data traffic stays local to the AP, and the controller is not able to filter out that broadcast traffic.
- **Convert Broadcast ARP Requests To Unicast** - If "Enabled", all broadcast ARP requests are converted to unicast and sent directly to the client. This configuration parameter is only intended for use for virtual APs in tunnel mode. In tunnel mode, all packets travel to the controller, so the controller is able to convert ARP requests directed to the broadcast address into unicast.


Assigning an Access Auth Profile

When you click the **Apply to Devices** button, the Access Auth Profile Assignments Screen appears. Click on the **Add/Remove Devices** button to bring up the Device Selection window. Select the switch(es)/wireless device(s) and click **OK**. The switch(es)/wireless device(s) will appear in the "Assign Switch" area.

- **Wired Switch** - To assign the profile to a wired switch (e.g., OmniSwitch Device), click on the switch, then click on the **Add Remove Ports** button to bring up the Port Selection window. Select the port(s) to which you want to assign the profile.
- **Wireless Device** - To assign the profile to a wireless device (e.g., OmniAccess Device), click on the device. Enter a VLAN to apply the Virtual AP Profile inside the Access Auth Profile. Each Virtual AP Profile can be assigned to multiple VLANs. Select an AP Group to assign to the Access Auth Profile. Only AP Groups that have been assigned to the selected device are displayed.

Click on the **Apply** button. The configuration will be applied and the assignment status displayed. Click **OK** to return to the Access Auth Profile Screen.



Editing an Access Auth Profile

Select the profile in the Access Auth Profile Screen and click on the Edit icon  to bring up the Edit Access Auth Profile Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server. (Note that you cannot edit the Access Auth Profile Name.) If the Access Auth Profile has been applied to any devices, you will have to re-apply the profile to those devices. You can also go to the Device Config - Access Auth Profile Screen to edit a profile on any device.

To "unassign" an Access Auth Profile from a device, go the Device Config - Access Auth Profile Screen and delete the profile from the device. To "unassign" a profile from specific device ports, go the Device Config - Access Auth Profile Screen and delete the profile from the device. Then return to the Access Auth Profile Screen, select the profile and re-assign it to the device, selecting only those ports to which you want to assign the profile.

For example, if you had assigned Access Auth Profile 1 to ports 1/1, 1/2, 1/3, and 1/4 on a device and you want to remove it from ports 1/3 and 1/4. You would go to the Device Config - Access Auth Profile Screen and delete Access Auth Profile 1 from the device. Then return to the Access Auth Profile Screen, select Access Auth Profile 1 and re-assign it to the device, selecting only ports 1/1 and 1/2 on the Device Selection Screen.

Deleting an Access Auth Profile

Select the profile in the Access Auth Profile Screen and click on the Delete icon , then click OK at the confirmation prompt. This removes the profile from the server. If the profile has been assigned to any devices, go to the Device Config - Access Auth Profile Screen to remove the profile from the device(s). Select the applicable device(s) in the Devices - Access Auth Profile Table, click on the Delete icon , then click **OK** at the confirmation prompt.

Wireless Profile Overview

The Unified Profile Wireless Profile feature is used to configure Virtual Access Point (AP) Profiles to enable authentication of wireless devices. APs advertise Wireless LANs (WLAN) to wireless clients by sending out beacons and probe responses that contain the WLAN's SSID and supported authentication and data rates. When a wireless client associates to an AP, it sends traffic to the AP's Basic Service Set Identifier (BSSID) which is usually the AP's MAC address. In a wireless network, an AP uses a unique BSSID for each WLAN. Thus, a physical AP can support multiple WLANs. The WLAN configuration applied to a BSSID on an AP is called a Virtual AP. You can configure and apply multiple virtual APs to an AP group or to an individual AP by defining one or more Virtual AP profiles.

The Wireless Profile feature enables you to configure multiple Virtual AP Profiles to provide different network access or services to users on the same physical network. For example, you could configure a WLAN to provide access to guest users, and another WLAN to provide access to employee users through


the same AP. You can configure the following authentication parameters for wireless devices and include these parameters in Access Authentication Profiles.

- SSID Profile - Used to configure SSID Profiles.
- AP Group - Used to configure AP Groups.

SSID Profile

The Unified Profile SSID Profile Screen displays all configured SSID Profiles and used to create, edit, assign, and delete SSID Profiles. An SSID Profile can be created and included in an Access Authentication Profile that can be assigned to wireless devices on the network.


Creating an SSID Authentication Profile

Click on the Create icon . Configure the Profile as described below, then click on the **Create** button.

SSID Authentication Profile

- **Profile Name** - User-configured profile name.
- **ESSID** - User configured name that uniquely identifies a wireless network (up to 32 characters). If the ESSID includes spaces, you must enclose it in quotation marks.
- **Hide SSID** - Enables/Disables the SSID name in beacon frames. Note that hiding the SSID does very little to increase security. (Default = Disabled)
- **Enable SSID** - Enables/Disables the SSID Profile.
- **Encryption** - The layer-2 authentication and encryption type used on this ESSID. DYNAMIC_WEP -
 - WEP with dynamic keys.
 - OPENSYSYSTEM - No authentication and encryption.
 - STATIC_WEP - WEP with static keys.
 - WPA_AES - WPA with AES encryption and dynamic keys using 802.1X.
 - WPA_PSK_AES - WPA with AES encryption using a preshared key.
 - WPA_PSK_TKIP - WPA with TKIP encryption using a preshared key.
 - WPA_TKIP - WPA with TKIP encryption and dynamic keys using 802.1X.
 - WPA2_AES - WPA2 with AES encryption and dynamic keys using 802.1X.
 - WPA2_PSK_AES - WPA2 with AES encryption using a preshared key.
 - WPA2_PSK_TKIP - WPA2 with TKIP encryption using a preshared key.
 - WPA2_TKIP - WPA2 with TKIP encryption and dynamic keys using 802.1X.


Editing an SSID Authentication Profile

Select the Profile in the SSID Authentication Profile Screen and click on the Edit icon  to bring up the Edit SSID Authentication Profile Screen. Edit the fields as described above then click on the **Save** button to save the changes to the server. Note that you cannot edit the profile name. .

Assigning an SSID Profile

When you click the **Apply To IAP Devices** button, the SSID Profile Assignments Screen appears. Click on the Add/Remove button to select device(s) and click the **Apply** button. The configuration will be applied and the assignment status displayed. Click **OK** to return to the SSID Profile Screen.

Deleting an SSID Authentication Profile


To delete a Profile(s), select the Profile(s) in the table and click on the Delete icon , then click **OK** at the confirmation prompt. If the profile is associated with an Access Authentication Profile, you will be presented


with warning prompt that you must remove the SSID Profile(s) from the Access Authentication Profile before it can be deleted. Remove the SSID Profile(s) from the Access Authentication Profile, and then delete the profile(s).

AP Group


The Unified Profile AP Group Screen displays all configured AP Groups and used to create, edit, assign, and delete AP Groups.

Creating an AP Group

Click on the Create icon . Configure the Group as described below, then click on the **Create** button.

- **Group Name-** User-configured name for the group.
- **Access Auth Profiles** - The Access Authentication Profile(s) associated with the group. Selecting an Access Auth Profile will allow for association of the AP Group to the correct Virtual AP Profile inside the Access Auth Profile. Select an Access Auth Profile from the drop-down menu or click on the Add icon  to go to the Access Auth Profile Screen and create a new one.


Editing an AP Group

Select the group in the AP Group Screen and click on the Edit icon  to bring up the Edit AP Group Screen. Edit the fields as described above then click on the **Save** button to save the changes to the server. Note that you cannot edit the group name.

Assigning an AP Group

You must first assign an AP Group to a controller, then associate the group with APs attached to that controller. First, select a group and click on the **Apply To Wireless Controllers** button. The AP Group Assignments Screen appears. Click on the **Add/Remove** button to select controller(s) and click the **Apply** button. The configuration will be applied and the assignment status displayed. Click **OK** to return to the AP Group Screen. Second, select an AP Group and click on the **Associate to APs** button. The Associate to APs Screen appears. Select APs and click on the **Apply** button.

Deleting an AP Group


To delete a group(s), select the group(s) in the table and click on the Delete icon , then click **OK** at the confirmation prompt.

AAA Server Profile


The Unified Profile AAA Server Profile Screen displays all configured AAA Server Profiles and is used to create, edit, and delete AAA Server Profiles for AOS 8.x Switches and Server Groups for Wireless Controllers. AAA Server Profiles are used to define specific AAA parameters that can be used in an Access Auth Profile or Captive Portal Profile.

Note: When an AAA Server Profile is assigned to a UNP Edge port/virtual AP through an Access Auth Profile, the parameter values defined in the profile will override any existing global AAA configuration for users authenticating on that port/virtual AP.


Creating an AAA Server Profile

Click on the Create icon . Enter a **Profile Name** and configure the Profile as described below, then click on the **Create** button.


Authentication Servers

- **802.1X Primary** - Select a Primary 802.1X Authentication Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the Add icon  to go to the Authentication Servers Application and create a new server. The Link takes you to the RADIUS Server Management Screen in the Authentication Server application. You can click on one of the other links on the left side of the screen to create a different Authentication Server type (LDAP, ACE, TACACS+).

For wireless devices, 802.1x Primary and Secondary Server configurations will help you to create 802.1x Authentication Server Group which will be used by Access Auth Profiles (Wireless AAA Server Profiles).

- **Captive Portal Primary** - Select a Primary Captive Portal Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the Add icon  to go to Authentication Servers Application and create a new Server. The Link takes you to the RADIUS Server Management Screen in the Authentication Server application. You can click on one of the other links on the left side of the screen to create a different Authentication Server type (LDAP, ACE, TACACS+).

Note: Captive Portal Primary and Secondary Server configurations are ignored for wireless devices.

- **MAC Primary**- Select a Primary MAC Authentication Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the Add icon  to go to Authentication Servers Application and create a new Server. The Link takes you to the RADIUS Server Management Screen in the Authentication Server application. You can click on one of the other links on the left side of the screen to create a different Authentication Server type (LDAP, ACE, TACACS+).

Note: For wireless devices, MAC Primary and Secondary Server configurations will help you to create a MAC Authentication Server Group that will be used by Access Auth Profiles (Wireless AAA Server Profiles). For IAP Devices, there is not a separate server for MAC Authentication. 802.1x Primary and Secondary Servers are used instead.

Accounting Servers

- **802.1X Primary** - Select a Primary 802.1X Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the "Add New" link to go to the RADIUS Server Management Screen and create a new
- Server. **Captive Portal Primary** - You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the "Add New" link to go to the RADIUS Server Management Screen and create a new Server.
- **MAC Primary** - Select a Primary MAC Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the "Add New" link to go to the RADIUS Server Management Screen and create a new Server.

Note: For wireless devices, Accounting Servers will help you to create an Accounting Radius Server Group that will be used in Access Auth Profiles (Wireless AAA Server Profiles). Captive Portal Primary and Secondary Servers are ignored. Wireless Devices

only accept Radius servers for Accounting. If you select another type, an error will occur when you try to apply the configuration to Wireless Controllers.

Advanced Settings

Advanced settings are not supported on wireless devices and will be ignored when applied to those devices.

MAC Auth

- **Session Timeout Trust Radius Status** - Enables/Disables the Session Timeout Trust Radius option for MAC Authenticated users. If Enabled, the switch will use the Session Timeout attribute received from the Authentication Server in an Accept-Accept message. If Disabled, the switch uses the locally configured timeout interval value (Default = Disabled).
- **Session Timeout Status** - Enables/Disables the Session Timeout option for MAC Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Session Timeout Interval. (Default = Disabled).
- **Session Timeout Interval** - The Session Timeout value, in seconds. When the Session Timeout value is reached, the authenticated users are logged out and the MAC address for each logged out user device is flushed. Note that when the Session Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again (Range = 12000 - 86400, Default = 43200).
- **Inactivity Timeout Status** - Enables/Disables the Inactivity Timeout option for MAC Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Inactivity Timeout Interval (Default = Disabled).
- **Inactivity Timeout Interval** - The Inactivity Timeout value, in seconds. Make sure the configured value is value greater than the MAC address aging time for the switch. If the Timeout Value is exceeded, the user is not logged out of the network if the MAC address aging time expires before the configured timeout value. Also note that when the Inactivity Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again.(Range = 60 - 1200, Default = 600)
- **Accounting Interim Trust Radius Status** - Enables/Disables the Accounting Interim Trust Radius option for MAC Authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again.
- **Accounting Interim Interval** - The amount of time between each interim accounting update for MAC accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Syslog Accounting Server IP Address** - The IP address of the Syslog Accounting Server.
- **Syslog Accounting Server UDP Port** - The port used to communicate with the Syslog Accounting Server (Default = 514).
- **Calling Station ID Type** - The RADIUS Calling Station ID attribute for MAC accounting sessions (**MAC** - sets the Calling Station ID to the MAC address of the user. **IP** - sets the Calling Station ID to the IP address of the user).

802.1X

- **Re-Authentication Timeout Trust Radius Status** - Enables/Disables the Session Timeout Trust Radius option for 802.1x Authenticated users. If Enabled, the Session-Timeout attribute value received from the RADIUS server overrides the locally configured value for the switch. (Default = Disabled).
- **Re-Authentication Timeout** - Enables/Disables the automatic re-authentication of authenticated 802.1X users (Default = Disabled).
- **Re-Authentication Interval** - The amount of time the switch waits, in seconds, before triggering re-authentication of 802.1X users. Note that when the re-authentication time interval is changed, the new value does not apply to existing authenticated 802.1X users until the user is flushed out or when the user is authenticated again. Any new 802.1X users are re-authenticated based on the current time interval setting. (Range = 600 - 7200, Default = 3600)
- **Accounting Interim Trust Radius Status** - Enables/Disables the Accounting Interim Trust Radius option for 802.1X authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again.
- **Accounting Interim Interval** - The amount of time between each interim accounting update for 802.1x accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Syslog Accounting Server IP Address** - The IP address of the Syslog Accounting Server.
- **Syslog Accounting Server UDP Port** - The port used to communicate with the Syslog Accounting Server (Default = 514).
- **Calling Station ID Type** - The RADIUS Calling Station ID attribute for MAC accounting sessions (**MAC** - sets the Calling Station ID to the MAC address of the user. **IP** - sets the Calling Station ID to the IP address of the user).

Captive Portal


- **Session Timeout Trust Radius Status** - Enables/Disables the Session Timeout Trust Radius option for Captive Portal Authenticated users. If Enabled, the switch will use the Session Timeout attribute received from the RADIUS server in an Accept-Accept message. If Disabled, the switch to use the locally configured timeout interval value (Default = Disabled).
- **Session Timeout Status** - Enables/Disables the Session Timeout option for Captive Portal Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Session Timeout Interval. (Default = Disabled).
- **Session Timeout Interval** - The Session Timeout value, in seconds. When the Session Timeout value is reached, the authenticated users are logged out and the MAC address for each logged out user device is flushed. Note that when the Session Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again (Range = 12000 - 86400, Default = 43200).
- **Inactivity Timeout Status** - Enables/Disables the Inactivity Timeout option for Captive Portal Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Inactivity Timeout Interval (Default = Disabled).
- **Inactivity Timeout Interval** - The Inactivity Timeout value, in seconds. Make sure the configured value is value greater than the MAC address aging time for the switch. If the Timeout Value is exceeded, the user is not logged out of the network if the MAC address aging time expires before the configured timeout value. Also note that when the Inactivity Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Range = 60 - 1200, Default - 600)

- **Accounting Interim Trust Radius Status** - Enables/Disables the Accounting Interim Trust Radius option for Captive Portal Authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again.
- **Accounting Interim Interval** - The amount of time between each interim accounting update for Captive Portal accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Syslog Accounting Server IP Address** - The IP address of the Syslog Accounting Server.
- **Syslog Accounting Server UDP Port** - The port used to communicate with the Syslog Accounting Server (Default = 514).
- **Calling Station ID Type** - The RADIUS Calling Station ID attribute for MAC accounting sessions (**MAC** - sets the Calling Station ID to the MAC address of the user. **IP** - sets the Calling Station ID to the IP address of the user).


RADIUS

- **NAS Port ID** - The RADIUS client NAS-Port attribute for authentication and accounting sessions. A text string (up to 31 characters) is used to define a NAS-Port identifier for the NAS-Port attribute. "Default" sets the NAS-Port attribute value to the chassis/slot/port of the user. The NAS-Port attribute value specified with this command is used in Account-Request messages and in Accounting-Request messages.
- **NAS ID** - The RADIUS client NAS-Identifier attribute for authentication and accounting sessions. A text string (up to 31 characters) is used to identify the switch (RADIUS client) in the NAS-Identifier attribute. "Default" sets the NAS-Identifier attribute to the system name of the switch. The NAS-Identifier attribute value specified with this command is used in both Account-Request and Accounting-Request messages.
- **Username Delimiter** - The delimiter character used to separate fields within a RADIUS Server User Name.
- **Password Delimiter** - The delimiter character used to separate fields within a RADIUS Server Password.
- **Calling Station Delimiter** - The delimiter character used to separate fields within a Calling Station ID.
- **Called Station Delimiter** - The delimiter character used to separate fields within a Called Station ID. **Username Case** - Indicates if the RADIUS Server User Name must be in Upper Case or Lower Case. **Password Case** - Indicates if the RADIUS Server Password must be in Upper Case or Lower Case.
- **Calling Station ID Case** - Indicates if the Calling Station ID must be in Upper Case or Lower Case.
- **Called Station ID Case** - Indicates if the Called Station ID must be in Upper Case or Lower Case.

Editing an AAA Server Profile

Select the profile in the AAA Server Profile Screen and click on the Edit icon  to bring up the Edit AAA Server Profile Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server. Note that if the AAA Server Profile has been applied to any devices through an Access Auth Profile or Captive Portal Profile, you will have to re-apply the associated Access Auth Profile or Captive Portal Profile to those devices to update the profile on the device(s).

Deleting an AAA Server Profile


Select the profile in the AAA Server Profile Screen and click on the Delete icon , then click **OK** at the confirmation prompt.

- If the profile has **not** been associated with an Access Auth Profile or Captive Portal Profile, the update will be applied and the status displayed. Click **OK** to return to the AAA Server Profile Screen.
- If the profile has been associated with an Access Auth Profile or Captive Portal Profile, the "Delete AAA Server Profile" confirmation prompt will appear listing any associated profiles. You must delete the AAA Server Profile from any associated profile(s) before returning to the AAA Server Profile Screen to delete the AAA Profile.

Access Role Profile

The Unified Profile Access Role Profile Screen displays all configured Access Role Profiles and is used to create, edit, and delete Access Role Profiles. An Access Role Profile contains the various UNP properties (e.g., QoS Policy List attached to the UNP, Captive Portal Authentication) for users assigned to the profile. In a wireless-centric network, an Access Role Profile is considered as a user role with which every client in the wireless-centric network is associated.

Creating an Access Role Profile

Click on the Create icon . Enter a **Profile Name** and configure the profile as described below, then click on the **Create** button. When you are finished, select the checkbox next to the profile and click on the **Apply to Devices** button to assign the profile to switches/wireless devices on the network.

Note: You can select a device type from the Highlight drop-down menu at the top of the screen to highlight configuration parameters for specific device types (6x, 7x, 8x).

Access Role Profile Attributes

- **Auth Flag** - Enables/Disables authentication (not supported on wireless devices and ignored when applied to those devices).
- **Mobile Tag Status** - Enables/Disables classification of tagged packets received on mobile ports (not supported on wireless devices and ignored when applied to those devices).
- **Redirect Status** - Enables/Disables Captive Portal Redirect (not supported on wireless devices and ignored when applied to those devices). Note that if Redirect Status is enabled, the Access Role Profile can only map to a VLAN when applying the profile to a device.
- **Policy List** - An Access Role Profile can also be configured with an existing Unified Policy List. The set of rules within the Unified Policy List are then applied to the traffic that passes through switches/wireless devices. Only one Unified Policy List is allowed per profile, but multiple profiles may use the same Policy List. Select a Unified Policy List for the profile from the drop-down menu. You can also click the "Add New" link to go to the **Unified Policy Lists** Screen to create a new one.
- **Captive Portal Auth** - Enables/Disables Captive Portal Authentication (not supported on wireless devices and ignored when applied to those devices).
- **Captive Portal Profile** - A Captive Portal Profile can be applied to AOS devices. Only one Captive Portal Profile is allowed per profile, but multiple profiles may use the same Captive Portal Profile. Select a Captive Portal Profile for the profile from the drop down menu. You can also click the "Add New" link to go to the **Captive Portal Profile** Screen to create a new one. (Not supported on wireless devices and ignored when applied to those devices.)
- **Inactivity Interval** - The amount of time, in seconds, before an authenticated device is

automatically logged out of the network due to inactivity (MAC address for the device has aged out). This timer value applies only to devices learned in the profile.

- **Max Ingress Bandwidth** - The maximum bandwidth limit allocated for ingress traffic on UNP ports assigned to the profile. If the maximum ingress bandwidth value is set to zero, all ingress traffic is allowed on the UNP port. (Not supported on AOS 7.3.4 switches and ignored when applied to those devices.)
- **Max Egress Bandwidth** - The maximum bandwidth limit allocated for egress traffic on UNP ports assigned to the profile. If the maximum egress bandwidth value is set to zero, all egress traffic is allowed on the UNP port. (Not supported on AOS 7.3.4 switches and ignored when applied to those devices.)

Max Ingress Depth or Max Default Depth (AOS 6) - The maximum ingress depth value that is applied to traffic on UNP ports that are assigned to the profile. This value determines how much the traffic can burst over the maximum ingress bandwidth rate. The maximum ingress depth value is configured in conjunction with the maximum ingress bandwidth parameter. When the ingress depth value is reached, the switch starts to drop packets. (Not supported on AOS 7.3.4 switches and ignored when applied to those devices.)

- **Max Egress Depth** - The maximum egress depth value that is applied to traffic on UNP ports that are assigned to profile. This value determines how much the traffic can burst over the maximum egress bandwidth rate. The maximum egress depth value is configured in conjunction with the maximum egress bandwidth parameter. When the egress depth value is reached, the switch starts to drop packets. (Not supported on AOS 7.3.4 switches and ignored when applied to those devices.)

ClearPass Attributes

- **Redirect URL** - The URL for the login page presented to a Guest user for Authentication using the Unified Access application (ClearPass). The URL is the path to the specific login page on the ClearPass Server that the user is directed to after initial authentication (e.g., Employee, Guest) (e.g., <http://clearpassuser.employee.page>). Note that this field is only available if a ClearPass Server is configured and connectivity has been established.

There are two (2) types of Enforcement Profiles created in ClearPass for each OmniVista Access Role Profile:

- **RADIUS CoA Enforcement Profile** - created in ClearPass with 'Radius_CoA' as the Type and 'CoA' as the Action. It will have at least two (2) RADIUS IETF attributes: Calling-Station-ID and Filter-ID. If a Redirect-URL is specified, the attribute called Alcatel-Redirection-URL is added to the URL in the OmniVista Access Role Profile.
- **Standard RADIUS Enforcement Profile** - created in ClearPass with 'Radius' as the Type and 'Accept' as the Action. If a Redirect-URL is specified, the attribute called Alcatel-Redirection-URL is added to the URL in the OmniVista Access Role Profile.

Wireless Settings

- **Upstream Bandwidth Contract kbit/s** - The maximum bandwidth for traffic from the switch to the client.
- **Downstream Bandwidth Contract kbit/s** - The maximum bandwidth for traffic from the client to the switch.

Assigning an Access Role Profile

When you click the **Apply To Devices** button, the Access Role Profile Assignments Screen appears. Select a VLAN Mapping Associate Method and/or Profile Policy, then assign the profile to switches/wireless devices on the network. Note that a VLAN must exist on a switch/wireless devices to configure VLAN Mapping.

Select Mapping Methods

You can map the Access Role Profile to a specific VLAN or service. Select a **Mapping Method**, then make a selection from the drop-down menu. Note that you can only use one mapping method for a profile.


- **Map to VLAN** - Maps the profile to a specific VLAN on network devices devices.
- **Map to SPB** - Maps the profile to an SPB Profile.
- **Map to VXLAN** - Maps the profile to a VXLAN Profile.
- **Map to Static Service** - Maps the profile to a Static Service.

Select Devices



After configuring the Mapping Method, select an option from the drop-down menu (Use Switch Picker/Use Topology) and click on the **Add/Remove Devices** button to select devices. The devices presented will vary according to your Mapping Method. For example, if you selected VLAN Number 3, only those devices on which VLAN 3 is configured would be displayed. After selecting devices, click on the **Apply** button to assign the Access Role Profile.

Note: You can also assign an Access Role Profile to a ClearPass Server. If a ClearPass Server is configured and connectivity established, the server will appear in the Device Selection Window in Blue.

Editing an Access Role Profile

Select the profile in the Access Role Profile List and click on the Edit icon  to bring up the Edit Access Role Profile Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server. (Note that you cannot edit the Access Auth Profile Name.) If the Access Role Profile has been applied to any devices, you will have to re-apply the profile to those devices. You can also go to the Device Config - Access Role Profile Screen to edit a profile on any device.

Deleting an Access Role Profile

Select the profile in the Access Role Profile Screen and click on the Delete icon , then click **OK** at the confirmation prompt. This removes the profile from the server. If the profile has been assigned to any devices, go to the Device Config - Access Role Profile Screen to remove the profile from the device(s). Select the applicable device(s) in the Devices - Access Role Profile Table, click on the Delete icon , then click **OK** at the confirmation prompt.

Access Classification

The Unified Profile Access Classification Screen displays all Access Classification Rules configured for Access Role Profiles and is used to create edit, and and delete Access Classification Rules (Access Classification Rules in AOS Switches and User Rules in wireless devices). Access Classification Rules are defined and associated with an Access Role Profile to provide an additional method for classifying a device into an Access Role Profile. If authentication is not available or does not return a profile name for whatever reason, Access Classification rules are applied to determine the profile assignment.

Creating an Access Classification Rule

Click on the Create icon **+**. Select a **Rule Type** from the drop-down menu. Configure the Rule as described below, select the **Access Role Profile** for which you want to configure the rule, then click on the **Create** button. When you are finished, click on the **Apply to Devices** button to assign the Rule to switches/ports on the network.


Access Classification Rules

- **MAC Rule (Both AOS and Wireless Devices)** - Defines a MAC Address Access Classification Rule for the specified UNP Access Role Profile. If the source MAC address of the device traffic matches the MAC address defined for the rule, the specified Access Role Profile is applied. Note that when a MAC Access Classification Rule is removed or modified, all MAC addresses classified with that rule are flushed.
 - **Name** - User-configured name for the MAC Rule.
 - **MAC Address** - The MAC address to be used for the rule. If the source MAC address of the device traffic matches the MAC address defined for the rule, the specified Access Role Profile is applied.
 - **VLAN Tag** - An optional VLAN Tag. If configured, traffic must also match this VLAN Tag in addition to the source MAC address.
 - **Customer Domain ID** - An optional Customer Domain ID to which this rule will apply. When a customer domain ID is configured for this rule, the rule is applied only to traffic received on UNP ports that are associated with the same domain ID. All UNP ports are automatically assigned to customer domain 0 at the time the port is configured as a UNP port.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **MAC Range Rule (AOS Devices only)** - Defines a MAC Address Range Access Classification Rule for the specified UNP Access Role Profile. If the source MAC address of the device traffic matches any of the MAC address within the range of MAC addresses, the specified profile is applied. Note that when a MAC Access Classification Rule is removed or modified, all MAC addresses classified with that rule are flushed.
 - **MAC Low Address** - MAC address that defines the low end of the range (e.g., 00:00:39:59:f1:00).
 - **MAC High Address** - MAC address that defines the high end of the range (e.g., 00:00:39:59:f1:90).
 - **VLAN Tag** - An optional VLAN Tag. If configured, traffic must also match this VLAN Tag in addition to the source MAC address.
 - **Customer Domain ID** - An optional Customer Domain ID to which this rule will apply. When a customer domain ID is configured for this rule, the rule is applied only to traffic received on UNP ports that are associated with the same domain ID. All UNP ports are automatically assigned to customer domain 0 at the time the port is configured as a UNP port.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **IP Address Rule (AOS Devices only)** - Defines an IP Address Access Classification Rule for the specified UNP Access Role Profile. If the source IP address of the device traffic matches the IP address defined for the rule, the specified Access Role Profile is applied.
 - **IP Network Address** - The IPv4 network address (e.g., 10.0.0.0, 171.15.0.0, 196.190.254.0). **IP Mask** - An IP address mask to identify the IP subnet for the interface (supports class-less masking).
 - **VLAN Tag** - An optional VLAN Tag. If configured, traffic must also match this VLAN Tag in addition to the source MAC address.
 - **Customer Domain ID** - An optional Customer Domain ID to which this rule will apply. When a customer domain ID is configured for this rule, the rule is applied only to traffic received on

UNP ports that are associated with the same domain ID. All UNP ports are automatically assigned to customer domain 0 at the time the port is configured as a UNP port.

- **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **VLAN Tag Rule** - Defines a VLAN Tag for the specified Access Classification Rule. If the source VLAN Tag of the device traffic matches the VLAN Tag defined for the rule, the specified Access Role Profile is applied.
 - **VLAN Tag** - The VLAN Tag used for the rule.
 - **Tag Position (7x only)** - The VLAN Tag position - Inner Tag (Default), Outer Tag.
 - **Customer Domain ID** - An optional Customer Domain ID to which this rule will apply. When a customer domain ID is configured for this rule, the rule is applied only to traffic received on UNP ports that are associated with the same domain ID. All UNP ports are automatically assigned to customer domain 0 at the time the port is configured as a UNP port.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **Location (Wireless Devices only)** - Defines a Location rule for the specified Access Role Profile. The specified Access Role Profile will be applied if the user location (AP name) matches with the value defined in the rule.
 - **Name** - The rule name.
 - **Location** - The AP location.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **ESSID (Wireless Devices only)** - Defines an Extended Service Set Identifier (ESSID) for the specified Access Role Profile. The specified Access Role Profile will be applied if the ESSID of AP (which client is associating) matches with the defined ESSID in the rule.
 - **Name** - The rule name.
 - **ESSID Value** - The ESSID of AP.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **DHCP Option (Wireless Devices only)** - Defines a DHCP signature ID rule for the specified Access Role Profile.
 - **Name** - The rule name.
 - **Signature ID** - The DHCP signature ID.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **DHCP Option 77 (Wireless Devices only)** - Defines a DHCP Option 77 rule for the specified Access Role Profile. The specified Access Role Profile will be applied if the user class identifier returned by DHCP server matches with the value defined in the rule.
 - **Name** - The rule name.
 - **Value** - The user class identifier returned by DHCP server.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **Encryption Type (Wireless Devices only)** - Defines an Encryption Type rule for the specified Access Role Profile. The specified Access Role Profile will be applied if the encryption type used by the client matches with the value defined in the rule.
 - **Name** - The rule name.
 - **Encryption Type** - The encryption type used by the client (e.g., WPA/WPA2 AES, Dynamic WEP).
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.

Editing an Access Classification Rule

Select the profile in the Classification Profile List and click on the Edit icon  to bring up the Edit Access Classification Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server. Note that if the Access Role Profile has been applied to any devices, you will have to re-apply the profile to those devices. You can also go to the Device Config - Access Classification Screen to edit a profile on any device.

Note: You cannot edit an Access Role Profile Name.

Assigning an Access Classification Rule

When you click the **Apply To Devices** button, the Access Classification Assignments Screen appears. Select a VLAN Mapping Associate Method and/or Profile Policy, then assign the profile to switches/wireless devices on the network. Note that a VLAN must exist on a switch/wireless devices to configure VLAN Mapping.

Select Mapping Methods

You can map the Access Classification Rule to a specific VLAN or service. Select a **Mapping Method**, then make a selection from the drop-down menu. Note that you can only use one mapping method for a profile.



- **Map to VLAN** - Maps the profile to a specific VLAN on network devices.
- **Map to SPB** - Maps the profile to an SPB Profile.
- **Map to VXLAN** - Maps the profile to a VXLAN Profile.

Select Devices

After configuring the Mapping Method, select an option from the drop-down menu (Use Switch Picker/Use Topology) and click on the **Add/Remove Devices** button to select devices. The devices presented will vary according to your Mapping Method. For example, if you selected VLAN Number 3, only those devices on which VLAN 3 is configured would be displayed. After selecting devices, click on the **Apply** button to assign the Access Classification Rule.

Note: During assignment of **Extended Rule**, port selection is offered, but it is optional. The rule can only be assigned to UNP Ports. Select **Port Range** and use the **From Port** and **To Port** options to assign the rule to a port(s). A port range can only be populated with consecutively higher numbered ports. Select Group Port to assign the rule to a port group, instead of a port(s).

Deleting an Access Classification Rule


To delete a rule(s), select the Rule(s) in the table and click on the Delete icon , then click **OK** at the confirmation prompt. This removes the profile from the server. If the profile has been assigned to any devices, go to the Device Config - Access Classification Screen to remove the profile from the device(s). Select the applicable device(s) in the Devices - Classification Profile List, click on the Delete icon , then click **OK** at the confirmation prompt.

Customer Domain

The Unified Profile Customer Domain Screen displays all configured Customer Domains create, edit, and delete Customer Domains. Customer Domains provide an additional method for segregating device traffic. A Customer Domain is identified by a numerical ID, which can be assigned to UNP ports and Access Classification Rules. By default, all UNP ports (bridge and access) and profile rules are assigned to the domain. The main benefits of Customer Domains is that they provide the ability to group physical UNP ports or link aggregates into one logical domain. Once a UNP port is assigned to a specific Customer Domain ID, only classification rules associated with the same ID are applied to that port.


An example of using customer domains would be to group UNP ports carrying traffic for a specific customer into the same domain (all Customer A ports assigned to Domain 2). Then assign VLAN and/or service profiles tailored for that customer to the same Domain ID (all profiles for Customer A assigned to Domain 2).

Creating a Customer Domain


Click on the Create icon  and complete the fields as described below. When you are finished, click on the **Create** button.

- **Customer Domain ID** - An ID number for the Customer Domain.
- **Customer Domain Description** - A description for the Customer Domain Profile.

Editing a Customer Domain

Select the profile in the Customer Domain List and click on the Edit icon  to bring up the Edit Customer Domain Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server. Note that if the Customer Domain Profile has been applied to any devices through an Access Classification Profile, you will have to re-apply the associated Access Classification Profile to those devices to update the profile on the device(s).

Deleting a Customer Domain


Select the profile in the Customer Domain List and click on the Delete icon , then click **OK** at the confirmation prompt. If the Customer Domain Profile has been applied to any devices through an Access Classification Profile, you will have to re-apply the associated Access Classification Profile to any devices to update the profile on the device(s).

- If the Customer Domain has **not** been associated with an Access Classification Profile, the update will be applied and the status displayed. Click **OK** to return to the AAA Server Profile Screen.
- If the Customer Domain has been associated with an Access Classification Profile, the "Delete" confirmation prompt will appear listing any associated profiles. You must delete the Customer Domain Profile from any associated profile(s) before returning to the Customer Domain Screen to delete the Customer Domain Profile.

SPB Profile

The Unified Profile SPB Profile Screen displays all configured Shortest Path Bridging (SPB) Profiles and is used to create, edit, and delete SPB Profiles. When you create an SPB Profile, you configure the parameters that can be mapped to an Access Role Profile. When a device is dynamically assigned to the profile through authentication or classification, and SPB Service Access Point (SAP) is automatically created using the specified profile parameters. Traffic from the device is then forwarded on the SAP.


Creating an SPB Profile

Click on the Create icon  and complete the fields as described below. When you are finished, click on the **Create** button.


- **SPB Profile Name** - The SPB Profile name.
- **Tag Value** - The VLAN tag information from classified traffic used to create the Service Access Point (SAP) for the traffic. If the traffic is untagged, the SAP is created with 0 as the encapsulation value (for example, 1/12:0).
- **ISID** - A service instance identifier (ISID) that is used to identify an SPB service in a provider backbone bridge (PBB) network. The valid range is 256 - 16777214.
- **BVLAN** - The VLAN ID number of an existing SPB backbone VLAN (BVLAN).

- **VLAN Translation** - Enables/Disables egress VLAN translation for the service.
- **Multicast Mode** - Select the multicast mode from the drop-down menu:
 - **Headend** - Specifies the head-end replication mode for the service.
 - **Tandem** - Specifies the tandem replication mode for the service.

Editing an SPB Profile

Select the profile in the SPB Profile List and click on the Edit icon  to bring up the Edit SPB Profile Screen. Edit the fields as described above then click on the **Apply** button. Note that you cannot edit the profile name.


Deleting an SPB Profile



Select the profile in the SPB Profile List and click on the Delete icon , then click **OK** at the confirmation prompt.

Far End IP


The Unified Profile Far End IP Screen displays all configured Far End IP Lists and is used to create, edit, and delete Far End IP Lists. Each IP address in a list is assigned to the Loopback0 interface of a far-end VXLAN node. The list name is assigned to an Access Role Profile through the mapping of VXLAN service parameters to the profile. This allows multiple far-end nodes to be associated with the service created for the VXLAN Network ID (VNID) specified in a VXLAN Profile.

Creating a Far End IP List


Click on the Create icon  and complete the fields as described below. When you are finished, click on the **Apply** button.

- **Name** - The Far End IP List name
- **IP Address** - Enter an IP address and click on the Add icon . Repeat to add additional IP addresses.
Click on the Delete icon  to remove an IP address.

Editing a Far End IP List

Select the list in the Far End IP Table and click on the Edit icon  to bring up the Edit Far End IP Screen. Edit the IP Address field as described above then click on the **Apply** button to save the changes. Note that you cannot edit the profile name.


Deleting a Far End IP List

Select the profile in the Far End IP Table and click on the Delete icon , then click **OK** at the confirmation prompt.

Static Service


Screen displays all configured Static Service mapping and is used to create, edit, and delete Static Service mapping. When you configure a Static Service, it is used to configure the mapping of an existing SPB or VXLAN service ID to an Access Role Profile. This type of profile mapping is only valid if the specified SPB or VXLAN service is already configured; the switch does not dynamically create the service. The specified service ID is then used to dynamically create a service access point (SAP) based on the specified tag value.

Creating a Static Service


Click on the Create icon  and complete the fields as described below. When you are finished, click on the **Add** button.

- **Name** - The SPB Profile name.
- **Tag Value** - The VLAN tag information from classified traffic used to create the Service Access Point (SAP) for the traffic. If the traffic is untagged, the SAP is created with 0 as the encapsulation value (for example, 1/12:0).
- **Service ID** - An existing (statically configured) numerical value that identifies a specific SPB or VXLAN service. The valid service ID range is 1–32767.

Editing a Static Service

Select the profile in the Static Service List and click on the Edit icon  to bring up the Edit Screen. Edit the fields as described above then click on the **Apply** button. Note that you cannot edit the profile name.


Deleting a Static Service

Select the profile in the Static Service List and click on the Delete icon , then click **OK** at the confirmation prompt.

VXLAN Profile


The Unified Profile VXLAN Profile Screen displays all configured VXLAN Profiles and is used to create, edit, and delete VXLAN Profiles. When you create a VXLAN, you configure the parameters that can be mapped to an Access Role Profile. When a device is dynamically assigned to the profile through authentication or classification, a VXLAN service access point (SAP) is automatically created using the specified profile parameters. Traffic from the device is then forwarded on the SAP.

Creating a VXLAN Profile


Click on the Create icon  and complete the fields as described below. When you are finished, click on the **Create** button.

- **Name** - The VXLAN Profile name.
- **Tag Value** - The VLAN tag information from classified traffic used to create the Service Access Point (SAP) for the traffic. If the traffic is untagged, the SAP is created with 0 as the encapsulation value (for example, 1/12:0).
- **VNID** - The VXLAN network identifier that identifies the VLAN segment from where the frames originate. This value is used to create the VXLAN service that is required to dynamically create the SAP.
- **Far End IP** - Select a Far End IP that contains the IP addresses for the far end VXLAN Tunnel End Points (VTEPs). The IP addresses in this list are used to dynamically create service distribution points (SDPs) for the VXLAN service.
- **Multicast IP Address** - The multicast IP address of the group to which this service will join.
- **VLAN Translation** - Enables/Disables egress VLAN translation for the service.
- **Multicast Mode** - Select the multicast mode from the drop-down menu:
 - **Headend** - Specifies the head-end replication mode for the service.
 - **Tandem** - Specifies the tandem replication mode for the service.
 - **Hybrid** - Specifies the hybrid replication mode for this service. This mode uses both the headend and tandem methods.

Editing a VXLAN Profile

Select the profile in the VXLAN Profile List and click on the Edit icon  to bring up the Edit VXLAN Profile Screen. Edit the fields as described above then click on the **Apply** button to save the changes. Note that you cannot edit the profile name.


Deleting a VXLAN Profile

Select the profile in the VXLAN Profile List and click on the Delete icon , then click **OK** at the confirmation prompt.

Global Configuration - AAA

The Unified Profile Global Configuration AAA Screen displays all configured Global AAA Profiles and used to create, edit, delete, and assign a Global AAA Profile. AAA Profiles are used to define specific AAA parameters that can be used in an Access Auth Profile or an Captive Portal Profile. This Global AAA Profile can be assigned and automatically applied to all UNP ports which have not been assigned an AAA Profile. In the absence of port template's AAA profile, the Global AAA Profile will be applied on AOS 8.x Switches.

Creating a Global AAA Server Profile

Click on the Create icon . Enter a **Profile Name** and configure the Profile as described below, then click on the **Create** button.

Authentication Servers

- **802.1X Primary** - Select a Primary 802.1X Authentication Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the "Add New" link to go to the Authentication Servers Application and create a new Server. (The Link takes you to the RADIUS Server Management Screen. You can click on one of the other links to create a different server type (LDAP, ACE, TACACS+).

Note: For wireless devices, 802.1x Primary and Secondary Server configurations will help you to create 802.1x Authentication Server Group which will be used by Access Auth Profiles (Wireless AAA Server Profiles).

- **Captive Portal Primary** - Select a Primary Captive Portal Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the "Add New" link to go to the Authentication Servers Application and create a new Server. (The Link takes you to the RADIUS Server Management Screen. You can click on one of the other links to create a different server type (LDAP, ACE, TACACS+).

Note: Captive Portal Primary and Secondary Server configurations are ignored for wireless devices.

- **MAC Primary**- Select a Primary MAC Authentication Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the "Add New" link to go to the Authentication Servers Application and create a new Server. (The Link takes you to the RADIUS Server Management Screen. You can click on one of the other links to create a different server type (LDAP, ACE, TACACS+).

Note: For wireless devices, MAC Primary and Secondary Server configurations will help you to create a MAC Authentication Server Group that will be used by Access Auth Profiles (Wireless AAA Server Profiles).

Accounting Servers

- **802.1X Primary** - Select a Primary 802.1X Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the click on the "Add New" link to go to the **RADIUS Server Management** Screen and create a new Server.
- **Captive Portal Primary** - Select a Primary Captive Portal Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the click on the "Add New" link to go to the **RADIUS Server Management** Screen and create a new Server.
- **MAC Primary** - Select a Primary MAC Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server. You can also click on the click on the "Add New" link to go to the **RADIUS Server Management** Screen and create a new Server.

Note: For wireless devices, Accounting Servers will help you to create an Accounting Radius Server Group that will be used in Access Auth Profiles (Wireless AAA Server Profiles). Captive Portal Primary and Secondary Servers are ignored. Wireless Devices only accept Radius servers for Accounting. If you select another type, an error will occur when you try to apply the configuration to Wireless Controllers.

Advanced Settings

Advanced settings are not supported on wireless devices and will be ignored when applied to those devices.

MAC Auth

- **Session Timeout Trust Radius Status** - Enables/Disables the Session Timeout Trust Radius option for MAC Authenticated users. If Enabled, the switch will use the Session Timeout attribute received from the Authentication Server in an Accept-Accept message. If Disabled, the switch uses the locally configured timeout interval value (Default = Disabled).
- **Session Timeout Status** - Enables/Disables the Session Timeout option for MAC Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Session Timeout Interval. (Default = Disabled).
- **Session Timeout Interval** - The Session Timeout value, in seconds. When the Session Timeout value is reached, the authenticated users are logged out and the MAC address for each logged out user device is flushed. Note that when the Session Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again (Range = 12000 - 86400, Default = 43200).
- **Inactivity Timeout Status** - Enables/Disables the Inactivity Timeout option for MAC Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Inactivity Timeout Interval (Default = Disabled).
- **Inactivity Timeout Interval** - The Inactivity Timeout value, in seconds. Make sure the configured value is value greater than the MAC address aging time for the switch. If the Timeout Value is exceeded, the user is not logged out of the network if the MAC address aging time expires before

the configured timeout value. Also note that when the Inactivity Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Range = 60 - 1200, Default - 600)

- **Accounting Interim Trust RADIUS Status** - Enables/Disables the Accounting Interim Trust Radius option for MAC Authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Default = Disabled)
- **Accounting Interim Interval** - The amount of time between each interim accounting update for MAC accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Calling Station ID Type** -The RADIUS Calling Station ID attribute for MAC accounting sessions (MAC - sets the Calling Station ID to the MAC address of the user. IP - sets the Calling Station ID to the IP address of the user).

802.1X

- **Re-Authentication Timeout Trust RADIUS Status** - Enables/Disables the Session Timeout Trust RADIUS option for 802.1x Authenticated users. If Enabled, the Session-Timeout attribute value received from the RADIUS server overrides the locally configured value for the switch. (Default = Disabled).
- **Re-Authentication Timeout Status** - Enables/Disables the automatic re-authentication of authenticated 802.1X users (Default = Disabled).
- **Re-Authentication Timeout Interval** - The amount of time the switch waits, in seconds, before triggering re-authentication of 802.1X users. Note that when the re-authentication time interval is changed, the new value does not apply to existing authenticated 802.1X users until the user is flushed out or when the user is authenticated again. Any new 802.1X users are re-authenticated based on the current time interval setting. (Range = 600 - 7200, Default = 3600)
- **Accounting Interim Trust RADIUS Status** - Enables/Disables the Accounting Interim Trust RADIUS option for MAC Authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Default = Disabled)
- **Accounting Interim Interval** - The amount of time between each interim accounting update for 802.1x accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Calling Station ID Type** -The RADIUS Calling Station ID attribute for MAC accounting sessions (MAC - sets the Calling Station ID to the MAC address of the user. IP - sets the Calling Station ID to the IP address of the user).

Captive Portal

- **Session Timeout Trust RADIUS Status** - Enables/Disables the Session Timeout Trust RADIUS option for Captive Portal Authenticated users. If Enabled, the switch will use the Session Timeout attribute received from the RADIUS server in an Accept-Accept message. If Disabled, the switch to use the locally configured timeout interval value (Default = Disabled).
- **Session Timeout Status** - Enables/Disables the Session Timeout option for Captive Portal Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Session Timeout Interval. (Default = Disabled).
- **Session Timeout Interval** - The Session Timeout value, in seconds. When the Session Timeout value is reached, the authenticated users are logged out and the MAC address for each logged out user device is flushed. Note that when the Session Timeout Interval is changed, the new value does


not apply to existing authenticated users until the user is flushed out or when the user is authenticated again (Range = 12000 - 86400, Default = 43200).

- **Inactivity Timeout Status** - Enables/Disables the Inactivity Timeout option for Captive Portal Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Inactivity Timeout Interval (Default = Disabled).
- **Inactivity Timeout Interval** - The Inactivity Timeout value, in seconds. Make sure the configured value is value greater than the MAC address aging time for the switch. If the Timeout Value is exceeded, the user is not logged out of the network if the MAC address aging time expires before the configured timeout value. Also note that when the Inactivity Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Range = 60 - 1200, Default - 600)
- **Accounting Interim Trust RADIUS Status** - Enables/Disables the Accounting Interim Trust RADIUS option for Captive Portal Authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Default = Disabled)
- **Accounting Interim Interval** - The amount of time between each interim accounting update for Captive Portal accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Calling Station ID Type** -The RADIUS Calling Station ID attribute for MAC accounting sessions (MAC - sets the Calling Station ID to the MAC address of the user. IP - sets the Calling Station ID to the IP address of the user).

RADIUS

- **NAS Port ID** - The RADIUS client NAS-Port attribute for authentication and accounting sessions. A text string (up to 31 characters) is used to define a NAS-Port identifier for the NAS-Port attribute. "Default" sets the NAS-Port attribute value to the chassis/slot/port of the user. The NAS-Port attribute value specified with this command is used in Account-Request messages and in Accounting-Request messages.
- **NAS ID** - The RADIUS client NAS-Identifier attribute for authentication and accounting sessions. A text string (up to 31 characters) is used to identify the switch (RADIUS client) in the NAS-Identifier attribute. "Default" sets the NAS-Identifier attribute to the system name of the switch. The NAS-Identifier attribute value specified with this command is used in both Account-Request and Accounting-Request messages.
- **Username Delimiter** - The delimiter character used to separate fields within a RADIUS Server User Name.
- **Password Delimiter** - The delimiter character used to separate fields within a RADIUS Server Password.
- **Calling Station Delimiter** - The delimiter character used to separate fields within a Calling Station ID.
- **Called Station Delimiter** - The delimiter character used to separate fields within a Called Station ID.
- **Username Case** - Indicates if the RADIUS Server User Name must be in Upper Case or Lower Case.
- **Password Case** - Indicates if the RADIUS Server Password must be in Upper Case or Lower Case.
- **Calling Station ID Case** - Indicates if the Calling Station ID must be in Upper Case or Lower Case.
- **Called Station ID Case** - Indicates if the Called Station ID must be in Upper Case or Lower Case.

Editing a Global AAA Server Profile


Select the profile in the AAA Server Profile Screen and click on the Edit icon  to bring up the Edit AAA Screen. Edit the fields as described above then click on the **Apply** button.

Note: You cannot edit the Profile Name.

Assigning a Global AAA Server Profile

When you click the **Apply To Devices** button, the AAA Assignment Screen appears. Click on the Add/Remove button to select the switch(es)/wireless device(s) and click the **Apply** button. The configuration will be applied and the assignment status displayed. Click **OK** to return to the AAA Screen. The **Override** button is used for most of the cases as **Apply** button. It is used to force assign the selected objects to the selected devices whether or not the objects exist on the device(s).

Deleting a Global AAA Server Profile

Select the Global AAA Profile in the AAA Screen and click on the Delete icon , then click **OK** at the confirmation prompt.

Device Config


Unified Profile Device Config Screens enable you to edit and delete Unified Profiles on specific network devices. When you select a profile on the left side of the screen (e.g., Access Auth Profile, Access Role Profile), all of the devices to which that profile type has been assigned are displayed. You can select a device and edit the profile parameters on the device, or select a device(s) and delete the profile from the device(s). The following screens are available:

- **Access Auth Profile** - Edit/Delete Access Authentication Profiles. An Access Auth Profile enables you to assign a pre-defined UNP port configuration to a port or linkagg, or specify them individually on each port to enable UNP port status and set the parameters for the authentication process for the port.
- **Wireless Profiles (SSID, AP Group, Virtual AP)** - Edit/Delete wireless profiles. Wireless Profiles can be created and included in an Access Authentication Profile that can be assigned to wireless devices on the network.
- **AAA Server Profile** - Edit/Delete AAA Server Profiles. AAA Server Profiles are used to define specific AAA parameters that can be used in an Access Auth Profile or Captive Portal Profile.
- **Access Role Profile** - Edit/Delete Access Role Profiles. An Access Role Profile contains the various UNP properties (e.g., QoS Policy List attached to the UNP, Captive Portal Authentication) for users assigned to the profile.
- **Access Classification** - Edit/Delete Access Classification Rules. Access Classification Rules are defined and associated with an Access Role Profile to provide an additional method for classifying a device into an Access Role Profile. If authentication is not available or does not return a profile name for whatever reason, Access Classification rules are applied to determine the profile assignment.
- **Far End IP** - Edit/Delete Far End IP Lists. Far End IP Lists allow multiple far-end nodes to be associated with the service created for the VXLAN Network ID (VNID) specified in a VXLAN Profile.
- **Diagnostics** - The Diagnostics Screen displays Unified Profile information for an end station which can be used to diagnose UNP Profile problems.
- **Global Configuration (AAA Profile)** - Edit/Delete Global AAA Profiles. AAA Profiles are used to define specific AAA parameters that can be used in an Access Auth Profile or an Captive Portal Profile.

Device Config - Access Auth Profile

The Unified Profile Device Config Access Auth Profile Screen displays information about all devices to which an Access Auth Profile has been assigned. You can edit the Access Auth Profile on a device, or delete the profile from a device(s). By default, all devices are displayed. However, you can display specific devices by selecting **Devices** from the **View By** drop-down menu and selecting specific devices.

Editing an Access Auth Profile

Select a device in the Access Auth Profile List and click on the Edit icon  to edit the field(s) as described below. When you are finished, click on the **Apply** button. Note that support for different parameters varies by device type. You can select an option from the "Highlight" drop-down menu at the top of the screen to highlight the parameters supported by specific devices (6x, 7x, 8x)

Default Settings

This section is used to configure basic settings for the profile.

- **AAA Server Profile** - The AAA Server Profile used to authenticate users on the port.
- **Port Bounce** - Enables/Disables Port Bounce. Always Enabled on wireless devices. This feature is required to handle scenarios where a client is switched from one VLAN to other after COA. If port bounce is enabled, the port will be administratively put down. This is to trigger DHCP renewal and re- authentication, if necessary.
- **MAC Auth** - Enables/Disables MAC Authentication for the port. Wireless devices do not contain this attribute in their configuration table. MAC Pass Alt attribute in the next section No Auth/Failure/Alternate is used for MAC Authentication on wireless devices.
- **802.1X Auth** - Enables/Disables 802.1X Authentication. Wireless devices do not contain this attribute in their configuration table. 802.1X Pass Alt attribute in the next section No Auth/Failure/Alternate is used for 802.1X Authentication on wireless devices.
- **Dynamic Service** - Select a dynamic mapping method, if applicable (SPB, VXLAN).
- **Customer Domain ID** - Select a Customer Domain ID for the profile, if applicable.

No Auth/Failure/Alternate

This section is used to configure the actions taken if a device assigned to the profile fails authentication.

- **Trust Tag** - Enables/Disables whether or not to trust the VLAN ID of a tagged packet to determine how the packet is classified. Enabling the trust VLAN ID tag option provides an implicit method of VLAN tag classification that will accept tagged traffic matching any of the existing UNPs without the need to create specific classification rules for those profiles.
- **Access Classification** - Enables/Disables device classification. Always Enabled on wireless devices (Default = Disabled).
- **Default Access Role Profile** - The Default Access Role Profile that users are assigned to after authentication. Note that for IAP devices the default Access Role Profile name must match the SSID Profile name in order for it to take effect.
- **802.1X Pass Alt** - The user shall be assigned a Pass-Alternate UNP in case the 802.1X authentication does not result in a valid UNP for the pass branch.
- **Bypass Status** - Enables/Disables 802.1X bypass. When 802.1X bypass is enabled, the user's 802.1X authentication method is skipped. The user enters directly mac-authentication or Access Classification based on the configuration on the UNP ports/Linkaggs. On wireless devices, this attribute corresponds to another attribute named I2-auth-fail-through, and this attribute must be combined with the MAC Allow EAP attribute to make I2-auth-fail-through attribute work (Default = Disabled).


- Bypass Status with ENABLED status combined with None MAC Allow EAP will disable 802.1X authentication, and I2-auth-fail-through is not ENABLED
- Bypass Status with ENABLED status combined with Fail MAC Allow EAP will enable I2-auth-fail-through.
- Other configurations of Bypass Status and MAC Allow EAP cause I2-auth-fail-through to be ignored on wireless devices.
- **Failure Policy** - The authentication method used if 802.1X authentication fails.
- **MAC Pass Alt** - The Access Role Profile the user is assigned to after passing authentication.
- **MAC Allow EAP** - Enables/Disables Extensible Authentication Protocol (EAP).

Advanced Settings

This section is used to configure advanced 802.1x authentication settings for the profile.

- **802.1X Tx Period** - Access Auth Profile 802.1x Tx period, in seconds.
- **802.1X Supp Timeout** - 802.1X Authentication Supp Timeout, in seconds.
- **802.1X Request** - 802.1X Authentication Max Request number.
- **Port Controlled Directions** - Configures whether network access control is applied to both incoming and outgoing traffic, or only applied to incoming traffic.

Deleting an Access Auth Profile

Select a profile(s) in the Access Auth Profile List and click on the Delete icon , then click **OK** at the confirmation prompt.

Device Config - SSID Profile

The Unified Profile Device Config SSID Profile Screen displays information about all devices to which an SSID Profile has been assigned. You can edit the SSID Profile on a device, or delete the it from a device(s). By default, all devices are displayed. However, you can display specific devices by selecting **Devices** from the **View By** drop-down menu an selecting specific devices.


Editing an SSID Profile

Select a device in the SSID Profile List and edit field(s) as described below. When you are finished, click on the **Apply** button.

- **Profile Name** - User-configured profile name.
- **ESSID** - User configured name that uniquely identifies a wireless network (up to 31 characters). If the ESSID includes spaces, you must enclose it in quotation marks.
- **Hide SSID** - Enables/Disables the SSID name in beacon frames. Note that hiding the SSID does very little to increase security. (Default = Disabled)
- **Enable SSID** - Enables/Disables the SSID Profile.
- **Encryption** - The layer-2 authentication and encryption type used on this ESSID.
 - DYNAMIC_WEP - WEP with dynamic keys.
 - OPENSYSYSTEM - No authentication and encryption.
 - STATIC_WEP - WEP with static keys.
 - WPA_AES - WPA with AES encryption and dynamic keys using 802.1X.
 - WPA_PSK_AES - WPA with AES encryption using a preshared key.
 - WPA_PSK_TKIP - WPA with TKIP encryption using a preshared key.
 - WPA_TKIP - WPA with TKIP encryption and dynamic keys using 802.1X.

- WPA2_AES - WPA2 with AES encryption and dynamic keys using 802.1X.
- WPA2_PSK_AES - WPA2 with AES encryption using a preshared key.
- WPA2_PSK_TKIP - WPA2 with TKIP encryption using a preshared key.
- WPA2_TKIP - WPA2 with TKIP encryption and dynamic keys using 802.1X.

Deleting an SSID Profile

Select a device(s) in the SSID Profile List and click on the Delete icon , then click **OK** at the confirmation prompt.

Device Config - AP Group


The Unified Profile Device Config AP Group Profile Screen displays information about all devices to which an AP Group Profile has been assigned. You can edit the AP Group Profile on a device, or delete the it from a device(s). By default, all devices are displayed. However, you can display specific devices by selecting **Devices** from the **View By** drop-down menu an selecting specific devices.

Editing an AP Group Profile

Select a device in the AP Group List and edit field(s) as described below. When you are finished, click on the **Apply** button.

- **Access Auth Profiles** - Select an Access Auth Profile from the drop-down list.

Deleting an AP Group Profile

Select a device(s) in the SSID Profile List and click on the Delete icon , then click **OK** at the confirmation prompt.

Device Config - Virtual AP

The Unified Profile Device Config Virtual AP Screen displays information about all devices to which a Virtual AP Profile has been assigned. You can edit the Virtual AP Profile on a device, or delete the it from a device(s). By default, all devices are displayed. However, you can display specific devices by selecting **Devices** from the **View By** drop-down menu an selecting specific devices.

Editing a Virtual AP Profile


Select a device in the Virtual AP List and edit field(s) as described below. When you are finished, click on the **Apply** button.

- **Access Auth Profile** - The Access Auth Profile associated with the Virtual AP Profile.
- **VLAN** - The VLAN the to which the profile is assigned.
- **SSID Profile** - The SSID Profile associated with the Virtual AP Profile.
- **Allowed Band** - The band(s) on which to use the Virtual AP:
 - **a** - 802.11a band only (5 GHz)
 - **g** - 802.11b/g band only (2.4 GHz)
 - **all** - Both 802.11a and 802.11b/g bands (5 GHz and 2.4 GHz). (Default).
- **Band Steering** - Enables/Disables Band Steering. Band Steering encourages dual-band capable clients to stay on the 5GHz band on dual-band APs. This frees up resources on the 2.4GHz band for single band clients like VoIP phones. The feature supports both campus APs and remote APs that have a virtual AP profile set to tunnel , split-tunnel or bridge forwarding mode. Note, however, that if a campus or remote APs have virtual AP profiles configured in bridge or split-tunnel forwarding mode but no virtual APs in tunnel mode, those APs will gather information about 5G-capable clients

independently and will not exchange this information with other APs that only have bridge or split-tunnel virtual APs configured.

- **Dynamic Multicast Optimization** - Enables/Disables Dynamic Multicast Optimization.
- **Dynamic Multicast Optimization Threshold** - The maximum number of high-throughput stations in a multicast group beyond which dynamic multicast optimization stops. (Range = 2 - 255, Default = 5)
- **Drop All Broadcast or Multicast Traffic** - If "Enabled", broadcast and multicast traffic is dropped. Do not enable this option for Virtual APs configured in bridge forwarding mode. This configuration parameter is only intended for use for Virtual APs in tunnel mode. In tunnel mode, all packets travel to the controller, so the controller is able to drop all broadcast traffic. When a Virtual AP is configured to use bridge forwarding mode, most data traffic stays local to the AP, and the controller is not able to filter out that broadcast traffic.
- **Convert Broadcast ARP Requests To Unicast** - If "Enabled", all broadcast ARP requests are converted to unicast and sent directly to the client. This configuration parameter is only intended for use for virtual APs in tunnel mode. In tunnel mode, all packets travel to the controller, so the controller is able to convert ARP requests directed to the broadcast address into unicast.
- **Forward Mode** - Controls whether data is tunneled to the controller using generic routing encapsulation (GRE), bridged into the local Ethernet LAN (for remote APs), or using a combination of both depending on the destination (e.g., corporate traffic goes to the controller, and Internet access remains local). All forwarding modes support band steering, TSPEC/TCLAS enforcement, 802.11k and station blacklisting.
 - **Tunnel** - The AP handles all 802.11 association requests and responses, but sends all 802.11 data packets, action frames, and EAPOL frames over a GRE tunnel to the controller for processing. The controller removes or adds the GRE headers, decrypts or encrypts 802.11 frames, and applies firewall rules to the user traffic as usual. Both remote and campus APs can be configured in tunnel mode.
 - **Bridge** - 802.11 frames are bridged into the local Ethernet LAN. When a remote AP or campus AP is in bridge mode, the AP (and not the controller) handles all 802.11 association requests and responses, encryption/decryption processes, and firewall enforcement. The 802.11e and 802.11k action frames are also processed by the AP, which then sends out responses as needed.
 - **Split Tunnel** - 802.11 frames are either tunneled or bridged, depending on the destination (e.g., corporate traffic goes to the controller, and Internet access remains local).
 - **Decrypt Tunnel** - Both remote and campus APs can be configured in decrypt-tunnel mode. When an AP uses decrypt-tunnel forwarding mode, that AP decrypts and decapsulates all 802.11 frames from a client and sends the 802.3 frames through the GRE tunnel to the controller, which then applies firewall policies to the user traffic.
- **Steering Mode** - Band steering supports the following three band steering modes.
 - **Force-5GHz** - The AP will try to force 5Ghz-capable APs to use that radio band.
 - **Prefer-5GHz** -The AP will try to steer the client to 5G band (if the client is 5G capable) but will let the client connect on the 2.4G band if the client persists in 2.4G association attempts.
 - (Default) **Band Balancing** - The AP tries to balance the clients across the two radios in order to best utilize the available 2.4G bandwidth. This feature takes into account the fact that the 5Ghz band has more channels than the 2.4 GHz band, and that the 5Ghz channels operate in 40MHz while the 2.5Ghz band operates in 20MHz.
- **Virtual AP Enable** - Enables/Disables the Wireless Authentication Profile.

Deleting a Virtual AP Profile

Select a device(s) in the Virtual AP List and click on the Delete icon , then click **OK** at the confirmation prompt.

Device Config - AAA Server Profile

The Unified Profile Device Config AAA Server Profile Screen displays information about all devices to which an AAA Server Profile has been assigned. You can edit the AAA Server Profile on a device, or delete the profile from a device(s). By default, all devices are displayed. However, you can display specific devices by selecting **Devices** from the **View By** drop-down menu and selecting specific devices.

Editing a AAA Server Profile

Select a device in the AAA Server Profile List and click on the Edit icon  to edit the field(s) as described below. When you are finished, click on the **Apply** button.

Authentication Servers

- **802.1X Primary** - Select a Primary 802.1X Authentication Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.

For wireless devices, 802.1x Primary and Secondary Server configurations will help you to create 802.1x Authentication Server Group which will be used by Access Auth Profiles (Wireless AAA Server Profiles).

- **Captive Portal Primary** - Select a Primary Captive Portal Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.

Note: Captive Portal Primary and Secondary Server configurations are ignored for wireless devices.

- **MAC Primary**- Select a Primary MAC Authentication Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.

Note: For wireless devices, MAC Primary and Secondary Server configurations will help you to create a MAC Authentication Server Group that will be used by Access Auth Profiles (Wireless AAA Server Profiles). For IAP Devices, there is not a separate server for MAC Authentication. 802.1x Primary and Secondary Servers are used instead.

Accounting Servers

- **802.1X Primary** - Select a Primary 802.1X Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.
- **Captive Portal Primary** - Select a Primary Captive Portal Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.
- **MAC Primary** - Select a Primary MAC Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.

Note: For wireless devices, Accounting Servers will help you to create an Accounting Radius Server Group that will be used in Access Auth Profiles (Wireless AAA Server Profiles). Captive Portal Primary and Secondary Servers are ignored. Wireless Devices only accept Radius servers for Accounting. If you select another type, an error will occur when you try to apply the configuration to Wireless Controllers.

Advanced Settings

Advanced settings are not supported on wireless devices and will be ignored when applied to those devices.

MAC Auth

- **Session Timeout Trust Radius Status** - Enables/Disables the Session Timeout Trust Radius option for MAC Authenticated users. If Enabled, the switch will use the Session Timeout attribute received from the Authentication Server in an Accept-Accept message. If Disabled, the switch uses the locally configured timeout interval value (Default = Disabled).
- **Session Timeout Status** - Enables/Disables the Session Timeout option for MAC Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Session Timeout Interval. (Default = Disabled).
- **Session Timeout Interval** - The Session Timeout value, in seconds. When the Session Timeout value is reached, the authenticated users are logged out and the MAC address for each logged out user device is flushed. Note that when the Session Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again (Range = 12000 - 86400, Default = 43200).
- **Inactivity Timeout Status** - Enables/Disables the Inactivity Timeout option for MAC Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Inactivity Timeout Interval (Default = Disabled).
- **Inactivity Timeout Interval** - The Inactivity Timeout value, in seconds. Make sure the configured value is value greater than the MAC address aging time for the switch. If the Timeout Value is exceeded, the user is not logged out of the network if the MAC address aging time expires before the configured timeout value. Also note that when the Inactivity Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again.(Range = 60 - 1200, Default = 600)
- **Accounting Interim Trust Radius Status** - Enables/Disables the Accounting Interim Trust Radius option for MAC Authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again.
- **Accounting Interim Interval** - The amount of time between each interim accounting update for MAC accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Syslog Accounting Server IP Address** - The IP address of the Syslog Accounting Server.
- **Syslog Accounting Server UDP Port** - The port used to communicate with the Syslog Accounting Server (Default = 514).
- **Calling Station ID Type** - The RADIUS Calling Station ID attribute for MAC accounting sessions (**MAC** - sets the Calling Station ID to the MAC address of the user. **IP** - sets the Calling Station ID to the IP address of the user).

802.1X

- **Re-Authentication Timeout Trust Radius Status** - Enables/Disables the Session Timeout Trust Radius option for 802.1x Authenticated users. If Enabled, the Session-Timeout attribute value received from the RADIUS server overrides the locally configured value for the switch. (Default = Disabled).
- **Re-Authentication Timeout** - Enables/Disables the automatic re-authentication of authenticated 802.1X users (Default = Disabled).

- **Re-Authentication Interval** - The amount of time the switch waits, in seconds, before triggering re-authentication of 802.1X users. Note that when the re-authentication time interval is changed, the new value does not apply to existing authenticated 802.1X users until the user is flushed out or when the user is authenticated again. Any new 802.1X users are re-authenticated based on the current time interval setting. (Range = 600 - 7200, Default = 3600)
- **Accounting Interim Trust Radius Status** - Enables/Disables the Accounting Interim Trust Radius option for 802.1X authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again.
- **Accounting Interim Interval** - The amount of time between each interim accounting update for 802.1x accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Syslog Accounting Server IP Address** - The IP address of the Syslog Accounting Server.
- **Syslog Accounting Server UDP Port** - The port used to communicate with the Syslog Accounting Server (Default = 514).
- **Calling Station ID Type** - The RADIUS Calling Station ID attribute for MAC accounting sessions (**MAC** - sets the Calling Station ID to the MAC address of the user. **IP** - sets the Calling Station ID to the IP address of the user).

Captive Portal

- **Session Timeout Trust Radius Status** - Enables/Disables the Session Timeout Trust Radius option for Captive Portal Authenticated users. If Enabled, the switch will use the Session Timeout attribute received from the RADIUS server in an Accept-Accept message. If Disabled, the switch to use the locally configured timeout interval value (Default = Disabled).
- **Session Timeout Status** - Enables/Disables the Session Timeout option for Captive Portal Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Session Timeout Interval. (Default = Disabled).
- **Session Timeout Interval** - The Session Timeout value, in seconds. When the Session Timeout value is reached, the authenticated users are logged out and the MAC address for each logged out user device is flushed. Note that when the Session Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again (Range = 12000 - 86400, Default = 43200).
- **Inactivity Timeout Status** - Enables/Disables the Inactivity Timeout option for Captive Portal Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Inactivity Timeout Interval (Default = Disabled).
- **Inactivity Timeout Interval** - The Inactivity Timeout value, in seconds. Make sure the configured value is value greater than the MAC address aging time for the switch. If the Timeout Value is exceeded, the user is not logged out of the network if the MAC address aging time expires before the configured timeout value. Also note that when the Inactivity Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Range = 60 - 1200, Default - 600)
- **Accounting Interim Trust Radius Status** - Enables/Disables the Accounting Interim Trust Radius option for Captive Portal Authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again.
- **Accounting Interim Interval** - The amount of time between each interim accounting update for Captive Portal accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Syslog Accounting Server IP Address** - The IP address of the Syslog Accounting Server.
- **Syslog Accounting Server UDP Port** - The port used to communicate with the Syslog


Accounting Server (Default = 514).

- **Calling Station ID Type** - The RADIUS Calling Station ID attribute for MAC accounting sessions (**MAC** - sets the Calling Station ID to the MAC address of the user. **IP** - sets the Calling Station ID to the IP address of the user).

RADIUS

- **NAS Port ID** - The RADIUS client NAS-Port attribute for authentication and accounting sessions. A text string (up to 31 characters) is used to define a NAS-Port identifier for the NAS-Port attribute. "Default" sets the NAS-Port attribute value to the chassis/slot/port of the user. The NAS-Port attribute value specified with this command is used in Account-Request messages and in Accounting-Request messages.
- **NAS ID** - The RADIUS client NAS-Identifier attribute for authentication and accounting sessions. A text string (up to 31 characters) is used to identify the switch (RADIUS client) in the NAS-Identifier attribute. "Default" sets the NAS-Identifier attribute to the system name of the switch. The NAS-Identifier attribute value specified with this command is used in both Account-Request and Accounting-Request messages.
- **Username Delimiter** - The delimiter character used to separate fields within a RADIUS Server User Name.
- **Password Delimiter** - The delimiter character used to separate fields within a RADIUS Server Password.
- **Calling Station Delimiter** - The delimiter character used to separate fields within a Calling Station ID.
- **Called Station Delimiter** - The delimiter character used to separate fields within a Called Station ID.
- **Username Case** - Indicates if the RADIUS Server User Name must be in Upper Case or Lower Case.
- **Password Case** - Indicates if the RADIUS Server Password must be in Upper Case or Lower Case.
- **Calling Station ID Case** - Indicates if the Calling Station ID must be in Upper Case or Lower Case.
- **Called Station ID Case** - Indicates if the Called Station ID must be in Upper Case or Lower Case.


Deleting a AAA Server Profile

Select a device(s) in the AAA Server Profile List and click on the Delete icon , then click **OK** at the confirmation prompt.

Device Config - Access Role Profile

The Unified Profile Device Config Access Role Profile Screen displays information about all devices to which a Access Role Profile has been assigned. You can edit the Access Role Profile on a device, or delete the profile from a device(s). By default, all devices are displayed. However, you can display specific devices by selecting **Devices** from the **View By** drop-down menu and selecting specific devices.


Editing an Access Role Profile

Select a device in the Access Role Profile List and click on the Edit icon  to edit the field(s) as described below. When you are finished, click on the **Apply** button. Note that support for different parameters varies by device type. You can select an option from the "Highlight" drop-down menu at the top of the screen to highlight the parameters supported by specific devices (6x, 7x, 8x).

Access Role Profile Attributes

- **Policy List** - An Access Role Profile can also be configured with an existing Unified Policy List. The set of rules within the Unified Policy List are then applied to the traffic that passes through switches/wireless devices. Only one Unified Policy List is allowed per profile, but multiple profiles may use the same Policy List. Select a Unified Policy List for the profile from the drop-down menu.
- **Max Ingress Bandwidth** - The maximum bandwidth limit allocated for ingress traffic on UNP ports assigned to the profile. If the maximum ingress bandwidth value is set to zero, all ingress traffic is allowed on the UNP port. (Not supported on AOS 7.3.4 switches and ignored when applied to those devices.)
- **Max Egress Bandwidth** - The maximum bandwidth limit allocated for egress traffic on UNP ports assigned to the profile. If the maximum egress bandwidth value is set to zero, all egress traffic is allowed on the UNP port. (Not supported on AOS 7.3.4 switches and ignored when applied to those devices.)
- **Max Ingress Depth or Max Default Depth (AOS 6)** - The maximum ingress depth value that is applied to traffic on UNP ports that are assigned to the profile. This value determines how much the traffic can burst over the maximum ingress bandwidth rate. The maximum ingress depth value is configured in conjunction with the maximum ingress bandwidth parameter. When the ingress depth value is reached, the switch starts to drop packets. (Not supported on AOS 7.3.4 switches and ignored when applied to those devices.)

Deleting an Access Role Profile

Select a device(s) in the Access Role Profile List and click on the Delete icon , then click **OK** at the confirmation prompt.

Device Config - Access Classification

The Unified Profile Device Config Access Classification Screen displays information about all devices to which a Access Classification Profile has been assigned. You can edit the Access Classification Rule on a device, or delete the profile from a device(s). By default, all devices are displayed. However, you can display specific devices by selecting **Devices** from the **View By** drop-down menu and selecting specific devices.

Editing an Access Classification Profile

Select a device in the Access Classification Profile List and edit the field(s) as described below. When you are finished, click on the **Apply** button. Note that the parameters you can edit depend on the Access Classification Profile assigned to the device.

- **MAC Rule (Both AOS and Wireless Devices)** - Defines a MAC Address Access Classification Rule for the specified UNP Access Role Profile. If the source MAC address of the device traffic matches the MAC address defined for the rule, the specified Access Role Profile is applied. Note that when a MAC Access Classification Rule is removed or modified, all MAC addresses classified with that rule are flushed.
 - **Name** - User-configured name for the MAC Rule.
 - **MAC Address** - The MAC address to be used for the rule. If the source MAC address of the device traffic matches the MAC address defined for the rule, the specified Access Role Profile is applied.
 - **VLAN Tag** - An optional VLAN Tag. If configured, traffic must also match this VLAN Tag in addition to the source MAC address.
 - **Customer Domain ID** - An optional Customer Domain ID to which this rule will apply. When a customer domain ID is configured for this rule, the rule is applied only to traffic received on

UNP ports that are associated with the same domain ID. All UNP ports are automatically assigned to customer domain 0 at the time the port is configured as a UNP port.

- **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **MAC Range Rule (AOS Devices only)** - Defines a MAC Address Range Access Classification Rule for the specified UNP Access Role Profile. If the source MAC address of the device traffic matches any of the MAC address within the range of MAC addresses, the specified profile is applied. Note that when a MAC Access Classification Rule is removed or modified, all MAC addresses classified with that rule are flushed.
 - **MAC Low Address** - MAC address that defines the low end of the range (e.g., 00:00:39:59:f1:00).
 - **MAC High Address** - MAC address that defines the high end of the range (e.g., 00:00:39:59:f1:90).
 - **VLAN Tag** - An optional VLAN Tag. If configured, traffic must also match this VLAN Tag in addition to the source MAC address.
 - **Customer Domain ID** - An optional Customer Domain ID to which this rule will apply. When a customer domain ID is configured for this rule, the rule is applied only to traffic received on UNP ports that are associated with the same domain ID. All UNP ports are automatically assigned to customer domain 0 at the time the port is configured as a UNP port.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **IP Address Rule (AOS Devices only)** - Defines an IP Address Access Classification Rule for the specified UNP Access Role Profile. If the source IP address of the device traffic matches the IP address defined for the rule, the specified Access Role Profile is applied.
 - **IP Network Address** - The IPv4 network address (e.g., 10.0.0.0, 171.15.0.0, 196.190.254.0).
 - **IP Mask** - An IP address mask to identify the IP subnet for the interface (supports classless masking).
 - **VLAN Tag** - An optional VLAN Tag. If configured, traffic must also match this VLAN Tag in addition to the source MAC address.
 - **Customer Domain ID** - An optional Customer Domain ID to which this rule will apply. When a customer domain ID is configured for this rule, the rule is applied only to traffic received on UNP ports that are associated with the same domain ID. All UNP ports are automatically assigned to customer domain 0 at the time the port is configured as a UNP port.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **VLAN Tag Rule** - Defines a VLAN Tag for the specified Access Classification Rule. If the source VLAN Tag of the device traffic matches the VLAN Tag defined for the rule, the specified Access Role Profile is applied.
 - **VLAN Tag** - The VLAN Tag used for the rule.
 - **Tag Position (7x only)** - The VLAN Tag position - Inner Tag (Default), Outer Tag.
 - **Customer Domain ID** - An optional Customer Domain ID to which this rule will apply. When a customer domain ID is configured for this rule, the rule is applied only to traffic received on UNP ports that are associated with the same domain ID. All UNP ports are automatically assigned to customer domain 0 at the time the port is configured as a UNP port.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **Location (Wireless Devices only)** - Defines a Location rule for the specified Access Role Profile. The specified Access Role Profile will be applied if the user location (AP name) matches with the value defined in the rule.
 - **Name** - The rule name.
 - **Location** - The AP location.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.

- **ESSID (Wireless Devices only)** - Defines an Extended Service Set Identifier (ESSID) for the specified Access Role Profile. The specified Access Role Profile will be applied if the ESSID of AP (which client is associating) matches with the defined ESSID in the rule.
 - **Name** - The rule name.
 - **ESSID Value** - The ESSID of AP.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **DHCP Option (Wireless Devices only)** - Defines a DHCP signature ID rule for the specified Access Role Profile.
 - **Name** - The rule name.
 - **Signature ID** - The DHCP signature ID.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **DHCP Option 77 (Wireless Devices only)** - Defines a DHCP Option 77 rule for the specified Access Role Profile. The specified Access Role Profile will be applied if the user class identifier returned by DHCP server matches with the value defined in the rule.
 - **Name** - The rule name.
 - **Value** - The user class identifier returned by DHCP server.
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.
- **Encryption Type (Wireless Devices only)** - Defines an Encryption Type rule for the specified Access Role Profile. The specified Access Role Profile will be applied if the encryption type used by the client matches with the value defined in the rule.
 - **Name** - The rule name.
 - **Encryption Type** - The encryption type used by the client (e.g., WPA/WPA2 AES, Dynamic WEP).
 - **Access Role Profile** - Select the Access Role Profile to use for the rule.

Deleting an Access Classification Profile

Select a device(s) in the Access Classification Profile List and click on the Delete icon , then click **OK** at the confirmation prompt.

Device Config - Far End IP


The Unified Profile Device Config Far End IP Screen displays information about all devices to which a Far End IP List has been assigned. You can edit the Far End IP List on a device, or delete the it from a device(s). By default, all devices are displayed. However, you can display specific devices by selecting **Devices** from the **View By** drop-down menu an selecting specific devices.

Editing a Far End IP List

Select a device in the Far End IP List and edit the field(s) as described below. When you are finished, click on the **Apply** button.

- **IP Address** - Enter an IP address and click on the Add icon  to add an address. Repeat to add additional IP addresses. Click on the Delete icon  to remove an IP address.

Deleting a Far End IP List

Select a device(s) in the Far End IP List and click on the Delete icon , then click **OK** at the confirmation prompt.

Device Config - Diagnostic

The Unified Profile Device Config Diagnostics Screen displays Unified Profile information for an end station connected to UNP Ports which can be used to diagnose UNP Profile problems. You can view information for a device by IP Address or MAC Address by selecting the applicable **Search by** criteria, entering the address and clicking on the **Locate** button.

Diagnostic List

- **Device Address** - The IP address of the device.
- **Port** - The slot/port on which the device was learned.
- **MAC Address** - The MAC address of the device.
- **Access Timestamp** - The login timestamp of the device.
- **User Name** - The name used to authenticate the device.
- **IP Address** - The IP address from which the device is sending packets.
- **VLAN** - The device VLAN.
- **Classification Source** - The Classification policy under which the device was learned.
- **Authentication Type** - The authentication type used to authenticate the device (e.g., MAC).
- **Authentication Status** - The status of authentication:
 - Idle
 - In Progress
 - Authenticated Failed
 - Failed Timeout
 - Failed No Server
 - Failed No Resources
- **IP Address Type** - The user IP address type. Currently, only IPv4 is supported.
- **Auth Server IP Used**- The IP address of the Authentication Server used to authenticate the device.
- **Auth Server IP Type** - The Authentication Server IP address type. Currently, only IPv4 is supported.
- **UNP Used** - The UNP used to classify the device.
- **User Role** - The UNP used to classify the device role.
- **User Role Source** - The UNP user role source.
- **Auth Fail Reason** - The authentication failure reason.
- **Auth Fail Retry Count** - The authentication failure retry count (number of times re-authentication is attempted after an authentication failure).
- **Classif Profile Rule** - The Classification Policy from which the device was learned.
- **Rest Access Status** - The MAC VLAN user Authentication Server status.
- **Role Rule** - The Classification Policy Rule used to classify the device.
- **Loc Policy Status** - The Location Policy status (Not Applicable/Pass/Fail).
- **Time Policy Status** - The Time Policy status (Not Applicable/Pass/Fail).
- **Cap Portal Status** - The Captive Portal status (Not Applicable/Pass/Fail).
- **Auth Server Used** - The name of the Authentication Server name used for the latest authentication session of the device.
- **Server Message** - The RADIUS server message displayed to the user.
- **Redirection URL** - The Redirect Server URL.
- **UNP From Auth Server** - The UNP returned by the Authentication Server for the device.
- **QMR Status** - The QMR status (Enabled/Disabled).
- **MC LAG Learning** - The Multi-Chassis Link Aggregate status (Enabled/Disabled).
- **SIP Call Type** - The SIP Call Type for the device (Normal Call/Emergency Call/Not In Call).
- **SIP Media Type** - The SIP Media Type for the device (Other/Audio/Video/None).

Device Config - Global Configuration - AAA

The Unified Profile Device Config Global Configuration - AAA Screen displays information about all devices to which a Global Configuration - AAA Profile has been assigned. You can edit the profile on a device, or delete the it from a device(s). By default, all devices are displayed. However, you can display specific devices by selecting **Devices** from the **View By** drop-down menu an selecting specific devices.

Editing a Global AAA Profile

Edit the field(s) as described below. When you are finished, click on the **Apply** button. Note that the fields vary depending on the device selected.

Authentication Servers

- **802.1X Primary** - Select a Primary 802.1X Authentication Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.

Note: For wireless devices, 802.1x Primary and Secondary Server configurations will help you to create 802.1x Authentication Server Group which will be used by Access Auth Profiles (Wireless AAA Server Profiles).

- **Captive Portal Primary** - Select a Primary Captive Portal Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.

Note: Captive Portal Primary and Secondary Server configurations are ignored for wireless devices.

- **MAC Primary**- Select a Primary MAC Authentication Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.

Note: For wireless devices, MAC Primary and Secondary Server configurations will help you to create a MAC Authentication Server Group that will be used by Access Auth Profiles (Wireless AAA Server Profiles).

Accounting Servers

- **802.1X Primary** - Select a Primary 802.1X Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.
- **Captive Portal Primary** - Select a Primary Captive Portal Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.
- **MAC Primary** - Select a Primary MAC Accounting Server for the Profile. You can also select Secondary, Tertiary, and Quaternary Backups, however each must be a different server.

Note: For wireless devices, Accounting Servers will help you to create an Accounting Radius Server Group that will be used in Access Auth Profiles (Wireless AAA Server Profiles). Captive Portal Primary and Secondary Servers are ignored. Wireless Devices only accept Radius servers for Accounting. If you select another type, an error will occur when you try to apply the configuration to Wireless Controllers.

Advanced Settings

Advanced settings are not supported on wireless devices and will be ignored when applied to those devices.

MAC Auth

- **Session Timeout Trust Radius Status** - Enables/Disables the Session Timeout Trust Radius option for MAC Authenticated users. If Enabled, the switch will use the Session Timeout attribute received from the Authentication Server in an Accept-Accept message. If Disabled, the switch uses the locally configured timeout interval value (Default = Disabled).
- **Session Timeout Status** - Enables/Disables the Session Timeout option for MAC Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Session Timeout Interval. (Default = Disabled).
- **Session Timeout Interval** - The Session Timeout value, in seconds. When the Session Timeout value is reached, the authenticated users are logged out and the MAC address for each logged out user device is flushed. Note that when the Session Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again (Range = 12000 - 86400, Default = 43200).
- **Inactivity Timeout Status** - Enables/Disables the Inactivity Timeout option for MAC Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Inactivity Timeout Interval (Default = Disabled).
- **Inactivity Timeout Interval** - The Inactivity Timeout value, in seconds. Make sure the configured value is value greater than the MAC address aging time for the switch. If the Timeout Value is exceeded, the user is not logged out of the network if the MAC address aging time expires before the configured timeout value. Also note that when the Inactivity Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again.(Range = 60 - 1200, Default - 600)
- **Accounting Interim Trust RADIUS Status** - Enables/Disables the Accounting Interim Trust Radius option for MAC Authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Default = Disabled)
- **Accounting Interim Interval** - The amount of time between each interim accounting update for MAC accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Calling Station ID Type** - The RADIUS Calling Station ID attribute for MAC accounting sessions (MAC - sets the Calling Station ID to the MAC address of the user. IP - sets the Calling Station ID to the IP address of the user).

802.1X

- **Re-Authentication Timeout Trust RADIUS Status** - Enables/Disables the Session Timeout Trust RADIUS option for 802.1x Authenticated users. If Enabled, the Session-Timeout attribute value received from the RADIUS server overrides the locally configured value for the switch. (Default = Disabled).
- **Re-Authentication Timeout Status** - Enables/Disables the automatic re-authentication of authenticated 802.1X users (Default = Disabled).
- **Re-Authentication Timeout Interval** - The amount of time the switch waits, in seconds, before triggering re-authentication of 802.1X users. Note that when the re-authentication time interval is changed, the new value does not apply to existing authenticated 802.1X users until the user is

flushed out or when the user is authenticated again. Any new 802.1X users are re-authenticated based on the current time interval setting. (Range = 600 - 7200, Default = 3600)

- **Accounting Interim Trust RADIUS Status** - Enables/Disables the Accounting Interim Trust RADIUS option for MAC Authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Default = Disabled)
- **Accounting Interim Interval** - The amount of time between each interim accounting update for 802.1x accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Calling Station ID Type** - The RADIUS Calling Station ID attribute for MAC accounting sessions (MAC - sets the Calling Station ID to the MAC address of the user. IP - sets the Calling Station ID to the IP address of the user).


Captive Portal

- **Session Timeout Trust RADIUS Status** - Enables/Disables the Session Timeout Trust RADIUS option for Captive Portal Authenticated users. If Enabled, the switch will use the Session Timeout attribute received from the RADIUS server in an Accept-Accept message. If Disabled, the switch to use the locally configured timeout interval value (Default = Disabled).
- **Session Timeout Status** - Enables/Disables the Session Timeout option for Captive Portal Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Session Timeout Interval. (Default = Disabled).
- **Session Timeout Interval** - The Session Timeout value, in seconds. When the Session Timeout value is reached, the authenticated users are logged out and the MAC address for each logged out user device is flushed. Note that when the Session Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again (Range = 12000 - 86400, Default = 43200).
- **Inactivity Timeout Status** - Enables/Disables the Inactivity Timeout option for Captive Portal Authenticated users. If Enabled, the user is automatically logged out of the network based on the configured Inactivity Timeout Interval (Default = Disabled).
- **Inactivity Timeout Interval** - The Inactivity Timeout value, in seconds. Make sure the configured value is value greater than the MAC address aging time for the switch. If the Timeout Value is exceeded, the user is not logged out of the network if the MAC address aging time expires before the configured timeout value. Also note that when the Inactivity Timeout Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Range = 60 - 1200, Default - 600)
- **Accounting Interim Trust RADIUS Status** - Enables/Disables the Accounting Interim Trust RADIUS option for Captive Portal Authenticated users. If Enabled, the Accounting Interim value received from the RADIUS server overrides the locally configured value. Note that when the Accounting Interim Interval is changed, the new value does not apply to existing authenticated users until the user is flushed out or when the user is authenticated again. (Default = Disabled)
- **Accounting Interim Interval** - The amount of time between each interim accounting update for Captive Portal accounting sessions, in seconds. (Range = 60 - 1200, Default - 600)
- **Calling Station ID Type** - The RADIUS Calling Station ID attribute for MAC accounting sessions (MAC - sets the Calling Station ID to the MAC address of the user. IP - sets the Calling Station ID to the IP address of the user).

RADIUS

- **NAS Port ID** - The RADIUS client NAS-Port attribute for authentication and accounting sessions. A text string (up to 31 characters) is used to define a NAS-Port identifier for the NAS-Port attribute. "Default" sets the NAS-Port attribute value to the chassis/slot/port of the user. The NAS-Port attribute value specified with this command is used in Account-Request messages and in Accounting-Request messages.
- **NAS ID** - The RADIUS client NAS-Identifier attribute for authentication and accounting sessions. A text string (up to 31 characters) is used to identify the switch (RADIUS client) in the NAS-Identifier attribute. "Default" sets the NAS-Identifier attribute to the system name of the switch. The NAS-Identifier attribute value specified with this command is used in both Account-Request and Accounting-Request messages.
- **Username Delimiter** - The delimiter character used to separate fields within a RADIUS Server User Name.
- **Password Delimiter** - The delimiter character used to separate fields within a RADIUS Server Password.
- **Calling Station Delimiter** - The delimiter character used to separate fields within a Calling Station ID.
- **Called Station Delimiter** - The delimiter character used to separate fields within a Called Station ID.
- **Username Case** - Indicates if the RADIUS Server User Name must be in Upper Case or Lower Case.
- **Password Case** - Indicates if the RADIUS Server Password must be in Upper Case or Lower Case.
- **Calling Station ID Case** - Indicates if the Calling Station ID must be in Upper Case or Lower Case.
- **Called Station ID Case** - Indicates if the Called Station ID must be in Upper Case or Lower Case.

Deleting a Global AAA Profile

Select a device(s) in the AAA Profile List and click on the Delete icon , then click **OK** at the confirmation prompt.

Profile Polling

The Unified Profile Profile Polling Screen is used to set the interval for polling devices the latest Unified Profile configurations. The current configured interval is displayed at the top of the screen. To change the interval, click on the "Reconfigure Poll Interval" link, set the new interval and click on the **Apply** button. (Range = 10 minutes to 24 hours, Default = 1 Hour).


You can also perform an immediate poll of devices. Select an option from the drop-down menu (Use Switch Picker/Use Topology) and click on the **Select Devices** button to choose the devices you want to poll. Click on the **Poll Now** button to poll the devices. When polling is complete, OmniVista will be updated with the latest Unified Profile information. Note that when polling is complete, you can click on the "Show More" link, then click on the "Details" link next to a device for detailed information on the polling operation.

Unified Policies

The PolicyView Unified Policies Screen application displays all configured Unified Policies and is used to create, edit, delete, and view Unified Policies. Unified Policies are QoS Policies that can be applied to both wireline and wireless devices. Unified Policies are created using a wizard that guides you through each of the steps needed to create the Policy and apply the Policy to devices in the network.

Note: Unified Policies are only displayed in the Unified Policies Table. They are **not** displayed with other configured QoS Policies in the Expert Mode Existing Policies Table.

Creating Unified Policy

Unified Policies are created using a wizard that guides you through each of the steps needed to create the policy and apply the policy to devices in the network. To create a Unified Policy, click on the Create icon . The wizard will then guide you through the following screens:

- Configuration - Basic policy configuration (e.g., Policy Name, Precedence)
- Device Selection - Specify the devices to which you will apply the policy
- Set Condition - Specify the conditions that must be true before traffic will be allowed to flow.
- Set Action - Specify parameters for the traffic that will flow.
- Validity Period - Specify the time period for the policy to be effective.
- Review - Review the policy details before creating the policy.

Note: As you configure a policy, conditions and actions are verified against the devices selected for the policy. If a condition or action is not supported by one of the selected devices, and error message will appear indicating the error and corrective action to be taken.


Applying a Unified Policy to the Network

After configuring and saving a policy(ies), you must apply the policy(ies) by notifying the switches in the network. When you click on the **Notify All** button, all of the policies listed in the Existing Unified Policies Table are applied to all of the devices configured for each policy. To apply the policy(ies) only to certain devices, select an option from the drop-down menu (Use Switch Picker/Use Topology), click on the **Select Device** button and select the device(s); then click the **Notify Selected** button.

After notifying the devices, you can view the status of the re-cache operation, by clicking on the **Status** button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the Audit application, including an indication of any error that may have occurred. Note that the re-cache operation for each switch occurs in a separate thread and may take some time. Any errors that occur will also be reported in the server.txt file, in the Audit application.

Note: When you notify network switches, all QoS-enabled switches flush their policy tables and reload policies from the LDAP repository, which is very expensive in terms of switch resources and time. It is recommended that you review all policies that you have created and apply them at the same time to minimize switch downtime.

Editing a Unified Policy

To edit a policy, select the policy in the Existing Unified Policies Table and click on the Edit icon . Use the wizard to make any edits. When you are done, apply the edited policy to the network.

Note that if you modify a policy and select different device types at the Device Selection Step (AOS/Wireless), a warning dialog will be displayed if the condition in the policy is not supported on one of the selected device types. For example: *“Condition mis-match: condition (L2 MACs and L4 Service) is not valid for selected device. Do you want to remove the mis-match conditions?”* If you select **Yes**, the mis-matched conditions will be removed from edited policy. Otherwise, the newly selected devices will be removed from Device Selection list.

Deleting Unified Policy

To delete a policy(ies), select the policy(ies) in the Existing Unified Policies Table and click on the Delete icon then click **OK** at the confirmation prompt.

Policy Information

The Existing Unified Policies Table displays information for all configured Policies. You can also click on a policy to view detailed information about the Policy (e.g., Condition, Action).

- **Policy Name** - The name of the Policy.
- **Scope** - The scope of the Policy (e.g., IP Filtering, Provisioned).
- **Precedence** - The Precedence value of the Policy (0 - 65535).
- **Status** - Indicates whether or not the Policy has been saved to the LDAP Server.
- **Enable** - Indicates whether or not the Policy is enabled.
- **Save** - Indicates whether or not the rule will be recorded during a snapshot command.
- **Log Matches** - Indicates whether or not matches to this rule are logged in the QoS Log.
- **Reflexive** - Indicates whether or not the Policy is reflexive. Reflexive Policy Rules allow specific return connections that would normally be denied (Yes/No).
- **Default List** - Indicates whether or not the Policy is saved to the Default Policy List. By default, a Policy Rule is added to this list when it is created. A Policy Rule remains a member of the Default List even when it is subsequently assigned to additional Policy Lists.
- **SLA Policy Trap** - Indicates whether or not an SLA Policy Trap is configured for the policy.

Config for Policy

The Unified Policies Config for Policy Screen is used to configure basic Policy parameters. When you have completed all of the parameters, click the **Next** button at the bottom of the screen or click on Device Selection on the left side of the screen to move to the next step.

- **Name** - The Policy name.
- **Precedence** - The Policy precedence. By default, the precedence field is pre-filled with the lowest unused precedence value (Range = 0 - 65535).

Click on **Show Advanced Options** to display and configure the options below:

- **Default List** - Adds the rule to the QoS Default Policy List. Default is
- **No Enabled** - Enables the policy. Default is **Yes**
- **Save** - Marks the policy rule so that it may be captured as part of the switch configuration. Default is **Yes**.
- **Log Matches** - Configures the switch to log messages about specific flows coming into the switch that match this policy rule. Default is **Yes**
- **Send Trap** - Enables traps for the Policy. Default is **No**
- **Reflexive** - Enables support for the Reflexive for the policy. Reflexive policies allow specific return connections that would normally be denied. Default is **Ignore**.

Note: The Config for Policy Screen for Unified Policies is similar to Config for Policy Screen in Expert mode. However, Unified Policies created for Wireless Controllers will accept the "No Reflexive" option.

Device Selection

The Unified Policies Device Selection Screen is used to select the switches to which you want to apply the Policy. Select an option (Use Switch Picker/Use Topology), and select the device(s). Click on the **Next** button at the bottom of the screen or click on Set Condition on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen and add/delete devices.

Note: In Expert Mode, you can only select AOS Devices for Policy creation. However, you can select wireline and wireless devices when creating Unified Policies. Also note that you

cannot select IAP Devices when creating Unified Policies.

Set Condition

The Unified Policies Set Condition Screen contains a list of Conditions that you can configure for the Policy (e.g., MAC Condition, IP Condition). When you create a Condition, the Condition(s) you configure must be true before traffic is allowed to flow. Click on a Condition to display the configuration options for the Condition. (Click again on the Condition to close the configuration options.) When you have completed all of the parameters for the Condition(s), click the **Next** button at the bottom of the screen or click on Set Action on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.

Conditions


A brief description of each Condition is provided below. Click the hyperlink for each Condition for detailed configuration instructions.

- **L2 MACs** - Create a Condition that applies the policy to traffic originating from a MAC address/group/range or to traffic flowing to a MAC address/group. (Note that for Wireless Controllers, MAC Addresses cannot contain wildcard characters).
- **L3 IPs** - Create a Condition that applies the policy to traffic originating from an IP address/network group or to traffic flowing to an IP address/network group. (Note that any IP address can be masked).
- **L3 DSCP/TOS** - Create a Condition that applies the policy to traffic with a specified value in either the DSCP (Differentiated Services Code Point) byte or in the IP TOS (IP Type of Service) byte. Both DSCP and IP TOS are mechanisms used to convey QoS information in the IP header of frames.
- **L4 Services** - Create a Condition that applies the policy to traffic flowing between two TCP or UDP ports, or to all traffic originating from a TCP or UDP port, or to all traffic flowing to a TCP or UDP port. You can also create a Condition using an existing service/service group.

Note: AOS Devices support most of above Conditions. However, Wireless Controllers do not. Please refer to detailed notes of each condition below for supported conditions.

L2 MACs

A MAC Condition applies the Policy to traffic flowing from/to a MAC Address/Group. Note that Layer 2 Conditions (conditions that specify MAC Addresses) are "lost" when traffic passes through a router. For this reason, it may be advisable to specify other types of Conditions (such as a Layer 3 Condition, which specifies IP Addresses) when traffic is expected to travel more than one router hop.

Select the parameter(s) you want to configure by selecting the applicable checkbox. Click on Single to configure a single MAC Address or Group to configure a MAC Group, then enter a MAC address or select a MAC Group from the drop-down menu. (You can also click the Add icon  to go to the **Groups** application and create a new MAC Group.)

- **Source MAC Address/MAC Group** - Configuring a Source MAC Address/Group Condition restricts the policy to traffic that flows from this MAC Address/Group only. If you do not select this option, you are effectively stating that the Source MAC Address/Group traffic is not a criterion for the policy.
- **Destination MAC Address/MAC Group** - Configuring a Destination MAC Address/Group Condition restricts the policy to traffic that flows to this MAC Address/Group only. If you do not select this option, you are effectively stating that the Destination MAC Address/Group traffic is not a criterion for the policy.


- **Source MAC Range** - Configuring a Source MAC Range Condition restricts the policy to traffic that flows from this MAC Range only. If you do not select this option, you are effectively stating that the Source MAC Range traffic is not a criterion for the policy.

Notes:

- Conditions that specify both a source and a destination MAC address may be rejected by some switch platforms as invalid. However, if you wish to create policies for both source and destination traffic, you can create one policy for the source traffic and a second policy for the destination traffic.
- MAC addresses may contain the wildcard character *. However, one * character must be entered for each individual hex digit in the MAC address: for example, **00435C:*******, not **00435C:***.
- The following MAC address ranges are assigned to Alcatel-Lucent Enterprise voice devices and Alcatel-Lucent Enterprise IP phones. You can create Conditions specifying these address ranges using the MAC Address tab.
 - Voice Devices
 - 00809F3A0000 - 00809F3AFFFF
 - 00809F3B0000 - 00809F3BFFFF
 - 00809F3C0000 - 00809F3CFFFF
 - IP Phones
 - 00809F3D0000 - 00809F3DFFFF
 - Multi-Media Devices
 - 00809F3E0000 - 00809F3EFFFF
 - 00809F3F0000 - 00809F3FFFFF
- Source MAC Range is not supported on AOS Devices.
- Source MAC Group and Destination MAC Address/MAC Group are not supported on Wireless Controllers.
- MAC Conditions are not supported on IAP Devices.

L3 IPs

An IP Condition applies the Policy to traffic originating from, or flowing to, an IP Address/Network group. Any IP Address can be masked. Note that a Condition that specifies both a Source and Destination IP Address/Network Group will be rejected by the switch as invalid. However, if you wish to create policies for both Source and Destination traffic, you can create one policy for the Source traffic and a second policy for the Destination traffic.

Select the parameter(s) you want to configure by selecting the applicable checkbox. For Source/Destination IP Address, click on Single to configure a single IP Address (and Shorthand or Subnet Mask, if applicable), or click on Group to configure a Network Group, then enter an IP Address or select a Network Group from the drop-down menu. (You can also click the Add icon  to go to the Groups application and create a new Network Group.)

- **Fragment (not available for Wireless controllers)** - Select this checkbox to restrict the policy to TCP packet fragments.
- **Source IP Address/Network Group** - Configuring a Source IP Address/Network Group Condition restricts the policy to traffic that flows from this IP Address or Subnet Mask/Network Group only. If you do not select this option, you are effectively stating that the Source IP Address or Subnet Mask/Network Group traffic is not a criterion for the policy.
- **Destination IP Address/Network Group** - Configuring a Destination IP Address/Network Group Condition restricts the policy to traffic that flows to this IP Address/Network Group only. If you do not select this option, you are effectively stating that the Destination IP Address or Subnet

Mask/Network Group traffic is not a criterion for the policy.

- **Multicast IP Address Range (not available for Wireless Controllers)** - Configuring a Multicast IP Address/Group Condition restricts the policy to traffic that flows to this IP Multicast Address Group only. If you do not select this option, you are effectively stating that the Destination IP Multicast Address or Subnet Mask/Group traffic is not a criterion for the policy.

Notes:

- When configuring an IP Address Condition, you can also click either the **Shorthand Mask** or **Subnet Mask** button to configure a Subnet Mask. If you are using a Shorthand Mask, select a value from the Shorthand Mask drop-down list. If you are using a full Subnet Mask, enter the mask in the IP Subnet Mask field. Note that the * wildcard character is not allowed in IP addresses.
Short hand Mask and Group are ignored when applying Unified Policies to Wireless Controllers.
- Source Group, Destination Group and Multicast are not supported on Wireless Controllers.

Important Note: When creating an IP Condition for a **NAT** Action you must specify a Network Group in the Condition. NAT will only work when both the Condition and Action specify network groups. To create a "One-to-Many" Condition and action, create a Network Group with a single entry for the Condition.

L3 DSCP/TOS

A DSCP/TOS Condition applies the Policy to incoming traffic that has a specified value in either the DSCP (Differentiated Services Code Point) byte or in the TOS (Type of Service) byte. Both DSCP and TOS are mechanisms used to convey QoS information in the IP header of frames. DSCP and TOS are mutually exclusive - you can use either DSCP or TOS but not both. Click on the applicable button (DSCP or TOS) and enter a value.




- **DSCP** - Defines the QoS treatment a frame is to receive from each network device. This is referred to as per-hop behavior. If you are using DSCP, you can define any value in the range 0 - 63 as the DSCP value in the IP header of the frame. Traffic that contains this value will match this condition.
- **TOS** - A TOS value creates a condition that applies the policy to traffic that has the specified TOS value in the IP header of frames. Enter any value from 0 - 7 to specify the value of the precedence field in the TOS byte that will match this condition. A value of 7 has the highest precedence and a value of 0 has the lowest.

Notes:

- Please refer to the Switch Release Notes for information on the specific QoS functions available on various current platforms and combinations of hardware/firmware.
- You cannot create a policy condition based on DSCP or TOS values for Wireless Controllers/IAPs. DSCP/TOS conditions are ignored when applying Unified Policies to Wireless Controllers/IAPs.

L4 Services

A Service Condition applies the policy to Service Protocol traffic (TCP or UDP) flowing from/to two TCP or UDP ports, or to traffic flowing from/to a TCP or UDP Service or Service Group. Select a type of Service Condition you want to configure, then configure the parameter(s) as described below.

- **Protocol Only** - Select **TCP** or **UDP** to create a condition for a Service Protocol only.
- **Port(s)** - To configure the Condition for a specific Service Port, select a **Source** and **Destination** Port from the drop-down menu to specify a specific port for the service you selected. You can also click on the Add icon  to go to the Groups application and create new Service Ports.
- **Service** - Select a Service from the drop-down menu. You can also click on the Add icon  to go to the Groups application and create a new Service.
- **Service Group** - Select a Service Group from the drop-down menu. You can also click on the Add icon  to go to the Groups application and create a new Service Group.

Notes:

- Wireless Controllers do not have source and destination ports. They only contain a unique service port. Therefore, you cannot specify both Source and Destination port for Wireless Controllers.

Set Action

The Unified Policies Set Action Screen contains a list of Actions that you can configure for the Unified Policy (e.g., QoS, NAT). A Policy Action enables you to specify the treatment traffic is to receive when it flows. This includes the priority the traffic will receive, its minimum and maximum output rates, and the values to which specified bits in the frame headers will be set upon egress from the switch. When the Conditions specified by the Policy Condition are true, traffic will flow as specified by the Policy Action.

Click on an Action to display the configuration options for the Action. (Click again on the Action to close the Action.) When you have completed all of the parameters for the Action(s), click the **Next** button at the bottom of the screen or click on Validity Period on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.

Actions

A brief description of each Action is provided below. Click the hyperlink for each Action for detailed configuration instructions.

- **QoS** - Create an Action to specify QoS actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action. Quality of Service applies to Session Type for wireless devices. Quality of Service is not supported on IAP devices and is ignored when applied to those devices.
- **TCM** - Create an Action to specify Tri-Color Marking (TCM) actions to impose on traffic that meets the configured policy condition(s). TCM provides a mechanism for policing network traffic by limiting the rate at which traffic is sent or received on a switch interface. TCM meters traffic based on user-configured packet rates and burst sizes and "marks" the metered packets as green, yellow, or red based on whether the traffic meets the configured rates. This "color marking" determines the packet's precedence when congestion occurs. TCM is not supported on wireless devices and is ignored when applied to those devices.

QoS

The QoS Policy Action option enables you to specify QoS actions to impose on traffic that meets the configured policy condition(s). When the conditions specified by the policy are true, traffic will flow as specified by the policy action.

- **Disposition** - Set the Action to **Accept** or **Drop** traffic that meets the configured condition(s).
- **Quality of Service (QoS) Parameters** - Specify the QoS priority the traffic will receive if it meets the configured condition(s).
 - **Platinum** priority provides the highest quality of service (and maps to a firmware priority of 7).
 - **Gold** provides the next-highest quality of service (and maps to a firmware priority of 5).
 - **Silver** provides the next-highest quality of service (and maps to a firmware priority of 3).
 - **Bronze** provides the same quality of service as best effort (and maps to a firmware priority of 1). A separate egress queue is maintained in the hardware for traffic of each different priority.
- **Output Flow Setting** (not supported on IAP Devices and is ignored when applied to those devices)
 - **Max Output Rate (kbits/sec)** - Specify the maximum amount of traffic, in kilobits-per-second, which is guaranteed to be transmitted from the port. Even if no other traffic exists, the output will be limited to the rate specified here.
 - **Set Color of Packet** - Enables/Disables Three Color Marking (TCM) for output traffic flows. This parameter is not supported on wireless devices and is ignored when applied to those devices.
- **Output Mapping** (not supported on IAP Devices and is ignored when applied to those devices)
 - **802.1p Priority Level** - If you want outgoing packets tagged with an 802.1p priority level, set the **802.1p Priority Level** field to any value between 0 to 7 to specify the desired outgoing 802.1p priority for the traffic. A value of **7** indicates the highest priority and a value of 0 indicates the lowest priority. Note that for ports that are configured for 802.1q, this value is used in the 802.1q header and indicates the outgoing priority of the frame. When a frame is de-queued for transmission, it is assigned the priority of the queue and mapped to the outgoing 802.1p priority. This priority is combined with the VLAN group ID to create the 802.1p/q header for transmission.
Note that if traffic matches the criteria specified by the policy condition, but the outgoing port does not support 802.1p tagging, the policy action will fail.
 - **DSCP/TOS** - Enable/Disable DSCP/TOS Precedence. The TOS byte is defined in RFC 791. This byte contains two fields. The precedence field is the three high-order bits (0-2) and is used to indicate the priority for the frame. The type of service field (bits 3-6) defines the throughput, delay, reliability, or cost for the frame; however, in practice these bits are not used. If you enable the **TOS Precedence** radio button, set the associated field to any value from 0-7 to specify the value that will be inserted into the precedence field of the TOS byte upon egress from the switch. A value of 7 has the highest precedence and a value of 0 has the lowest precedence. Note that you can enable **either** the DSCP or the TOS Precedence radio button to specify the mechanism you want to use (if any) to convey QoS information in the IP header of frames. DSCP and TOS are mutually exclusive. You can use either DSCP or TOS, but not both.

TCM

The TCM Policy Action option enables you to specify Three-Color Marking (TCM) actions action to impose on traffic that meets the configured policy condition(s). TCM provides a mechanism for policing network traffic by limiting the rate at which traffic is sent or received on a switch interface. TCM meters traffic based on user- configured packet rates and burst sizes and "marks" the metered packets as green, yellow, or red based on whether the traffic meets the configured rates. This "color marking" determines the packet's precedence when congestion occurs. TCM is not supported on IAP Devices and is ignored when applied to those devices.

- **Committed Traffic Policing**
 - **Committed Information Rate** - The maximum amount of bandwidth, in kbits-per-second, for all traffic that ingresses on the port.

Validity Period

The Unified Policies Validity Period Screen enables you to add a validity period to a condition by specifying the time periods when the policy is active and enforced. Four pre-configured policy validity periods are provided in the drop-down list in the **Policy Validity Periods** pane. They are **AllTheTime**, **Weekdays**, **Weekends**, and **WorkingDay**. You can also create **Custom** validity periods that are enforced during a specific timeframe.

When you have completed all of the parameters, click the **Next** button at the bottom of the screen or click on Review on the left side of the screen to move to the next step. If necessary, you can also click the **Back** button to return to the screen.

Note: The pre-configured validity period **AllTheTime** is the default and is automatically assigned to the condition when the **Ignore Validity Period in defining Policy Condition** checkbox is checked. You can configure a validity period when configuring an IP Condition or Service Condition. If you do not specify an IP or Service Condition, the configured period is not applied for Wireless Controllers.

Advanced Wireless Settings

For Wireless Controllers, you can specify an absolute period or a periodic period.

- **Absolute**- Specifies an absolute time range, with a specific start and/or end time and date.
- **Periodic** - Specifies a recurring time range. Specify the start and end time and all days or selected days of the week.

Review

The Unified Policies Review Screen is used to review the Policy configuration before saving the Policy. After reviewing the Policy, click the **Create** button to save the policy to the LDAP Server. You can also click the **Back** button to return to a previous screen.

Unified Policy List

The PolicyView Unified Policy List Screen displays all configured Unified Policy Lists, including the Unified Policies included in each list, and is used to create, edit, delete, view and apply Unified Policy Lists. A Unified Policy List is a set of Unified Policies that are grouped together and assigned to devices as a group. A Unified Policy List can be applied to an AOS Switch or ClearPass Server. A Unified Policy List can be applied to wireless devices as part of an Access Role Profile. Access Role Profiles are configured in the Unified Access application (Unified Access - Device Config - Access Role Profile).

Unified Policy List Information


The following information is displayed for each Unified Policy contained in the Unified Policy List. (Click on a Unified Policy List to display the Unified Policies contained in the list.)

- **Name** - The name of the Unified Policy.
- **Condition** - The Unified Policy Condition information (e.g., IP Policy Condition would display the

Source/Destination/Multicast IP address of the condition).


- **Action** - The Unified Policy Action to take if the traffic matches the Policy condition (e.g., QoS Accept/Drop)
- **Precedence** - The Precedence value of the Unified Policy (0 - 65535).
- **Validity Period** - The configured validity period for the Unified Policy.

Creating a Unified Policy List

Click on the Create icon . The Create Unified Policy List Wizard appears. Complete the screens as described below, then click on the **Create** button.

Config for Policy List

Enter a **Name** for the Unified Policy List and select the Unified Policies you want to include in the list from the

Add Unified Policies drop-down menu. (All of the currently-configured Unified Policies appear in the list. You can also click the Add icon  to go to Unified Policies Screen and create a new Unified Policy(ies) to add to the list.) When you select a Unified Policy from the drop-down menu, the Unified Policy will appear in a table below, so you can review the Unified Policy and modify the Precedence value, if needed.

If you are assigning a Policy List to both wired and wireless devices, select an option from the drop-down menu at the bottom of the table to override the default behavior of the devices. The default behavior for traffic that does not match a policy is different for wired and wireless devices. For wired devices, the default behavior is to "accept" the traffic. For wireless devices, the default behavior is to "deny" the traffic. For example, if you create a Source IP Policy for a single IP address, by default wired devices would accept traffic that does not come from that IP address while wireless devices would drop the traffic. If you are assigning a Policy List to both wired and wireless devices, select an option from the drop-down menu to override device default behavior. All devices that the policy is assigned to will then follow this default behavior.

- **OV-L3-AcceptAllPolicy** - Traffic that does not match any of the policies will be accepted on all devices.
- **OV-L3-DenyAllPolicy** - Traffic that does not match any of the policies will be denied on all devices. **Device-Default** - Traffic that does not match any of the policies will be accepted/denied according to the device's default behavior. If you do not make a selection from the drop-down menu, this option is automatically used.

Review the Policy List configuration(s) in the table, then click the **Create** button. The new Unified Policy List will appear on the Unified Policy Lists Screen.

Note: The Wireless User Role contains a QoS rule and an Access List, which is a set of ACLs. For User Role, Wireless Controllers support two (2) QoS attributes - Bandwidth Contract - Upstream and Downstream. However, OmniVista only supports configuring Downstream Bandwidth. Additionally, the User Role can contain only a single Bandwidth Contract. So if the Unified Policy List contains more than one QoS Rule, OmniVista will display an error message: "Unified Policy List can't contain more than one QOS Action."

Device Selection

Select an option from the drop-down menu (Use Switch Picker/Use Topology), click on the **Select Devices** button and select the device(s) to which you want to apply the Policy List.

Applying a Unified Policy List to the Network



After configuring and saving a Unified Policy List, you must apply the list by notifying the devices in the network. A Unified Policy List can be assigned to AOS Switches and/or ClearPass Servers. When you click on the **Notify All** button, all of the configured Policy Lists are applied to all of the devices configured for each policy. To apply a list only to certain devices, select an option from the drop-down menu (Use Switch Picker/Use Topology), click on the **Select Devices** button and select the device(s); then click the **Notify Selected** button.

Note: A Unified Policy List can be applied to wireless devices as part of an Access Role Profile.


After notifying the devices, you can view the status of the re-cache operation, by clicking on the **Status** button to view the Devices Pending Notification Table. In addition, you can view the success or failure of the re-cache operation for each switch in the policy.log file of the **Audit** application, including an indication of any error that may have occurred. Note that the re-cache operation for each switch occurs in a separate thread and may take some time. Any errors that occur will also be reported in the server.txt file, in the **Audit** application.

Note: When you notify network switches, all QoS-enabled switches flush their policy tables and reload policies from the LDAP repository, which is very expensive in terms of switch resources and time. It is recommended that you review all policies that you have created and apply them at the same time to minimize switch downtime.

Editing a Unified Policy List

You can edit the Unified Policies included in a Unified Policy List or edit the Precedence value of any Unified Policy in the list. Select a Unified Policy List and click on the Edit icon . The Edit Unified Policy List Screen appears. Click on the **Add Unified Policies** drop-down menu . (All of the currently-configured Unified Policies appear in the list. You can also click the Add icon  to go to Unified Policies Screen and create a new Unified policy(ies) to add to the list.) Select/unselect Unified Policies to add/remove them from the Unified Policy List. When you are finished editing the Unified Policy, click the **Update** button. The updated Unified Policy List will appear on the Unified Policy Lists Screen.

Deleting a Unified Policy List

To delete a Unified Policy List(s), select the list(s), click on the Delete icon , then click **OK** at the confirmation prompt. Note that you cannot delete a Unified Policy List that is associated with an Access Role Profile. To delete the list, you must first remove it from associated Access Role Profile.

Unified Policy List Information

Click on a Unified Policy List to display the information about the Policies contained in the list. The following information is displayed for each Unified Policy contained in the Unified Policy List.

- **Policy Name** - The name of the Policy.
- **Action** - The Policy Action to take if the traffic matches the Policy condition (e.g., QoS)
- **Accept/Drop Condition** - The Policy Condition information (e.g., IP Policy Condition would display the Source/Destination/Multicast IP address of the condition).
- **Precedence** - The Precedence value of the Policy (0 - 65535).
- **Validity Period** - The configured validity period for the Policy.

mDNS

The Unified Access mDNS application is used to configure the Multicast Domain Name System (mDNS) protocol. mDNS is used by "Zero Configuration Networking" solutions such as Apple's Bonjour, Avahi LGPL, and Linux NSS-mDNS. mDNS is a resolution service that is used to discover services on a LAN. mDNS allows the resolution of host names to IP addresses within small networks without the need of a conventional DNS server. The mDNS protocol uses IP multicast User Datagram Protocol (UDP) packets and is implemented by Apple Bonjour, Avahi (LGPL), and Linux NSS-mDNS. In a BYOD network, mDNS is leveraged by providing wireless guests and visitors access to network devices, such as printers.

mDNS

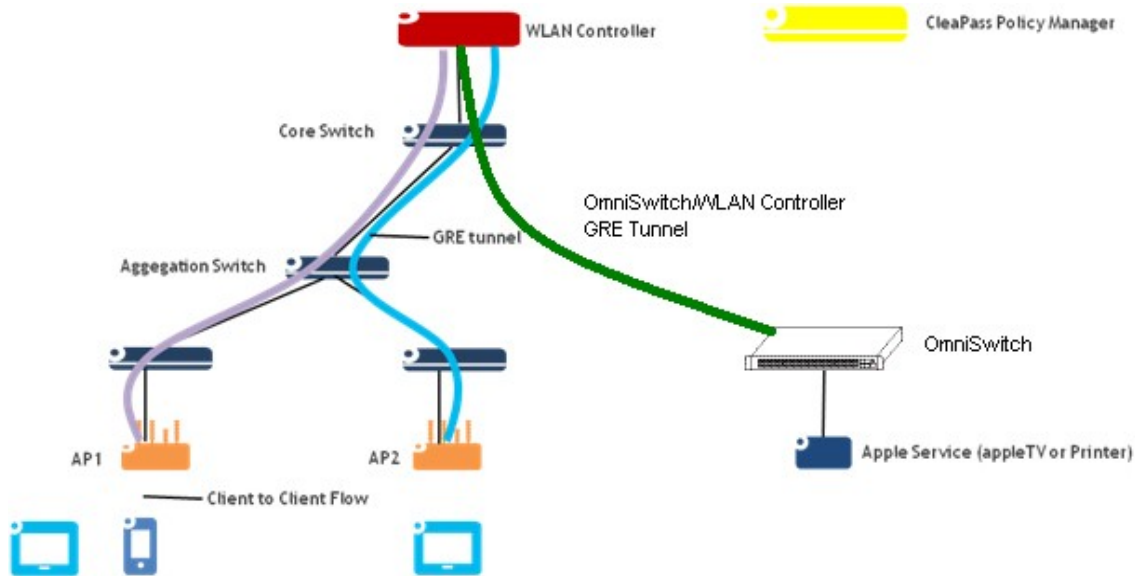
The screenshot displays the mDNS Configuration interface. At the top, there is a navigation bar with 'Home', 'Security', and 'mDNS'. Below this, the 'mDNS Configuration' section includes a search bar, a 'Reset' button, and options to download a '.csv' file, 'Add to Report', and 'Print'. The main content is a table with the following data:

Name	Device	mDNS GRE Tunnel	mDNS Admin Status	mDNS Oper Status
NMS-6850E-234	10.255.225.234	grr	Enabled	
shasta242	10.255.225.242	ag	Disabled	Down
vxTarget3	10.255.225.241	gretn1	Enabled	

At the bottom of the table, there is a 'Show: All' dropdown menu and a pagination indicator showing 'Showing All 3 rows'.

mDNS Flow

mDNS is configured on an OmniSwitch by creating a GRE Tunnel between the OmniSwitch and a Wireless LAN Controller. The following figure below provides a sample mDNS workflow setup. The wireless clients connected to Access point 1 (AP1) or Access Point 2 (AP2) request the mDNS service offered.



The mDNS feature is enabled on the OmniSwitch to support the mDNS service. A Layer 2 GRE tunnel interface is configured from the WLAN controller to the OmniSwitch to relay the mDNS messages. The mDNS message from the Bonjour capable wired service device is encapsulated and relayed from the OmniSwitch to the configured WLAN controller over the GRE tunnel. The WLAN controller then relays the mDNS messages received via the OmniSwitch GRE tunnel to the APs over the AP GRE tunnels.

Note that the WLAN controller uses a multicast optimization algorithm and forwards Bonjour response messages to targeted user devices, instead of all devices on all APs. This limits the unnecessary flooding of the Bonjour/mDNS traffic to improve the Wi-Fi performance.

mDNS Configuration Screen

The mDNS Configuration Screen is used to create, edit, delete, and view mDNS configuration on your network. Only **one** mDNS configuration per device is supported.

Configuring mDNS for a Device


Before you begin, configure the GRE tunnel interface before attempting to associate the interface with the mDNS tunnel relay. The GRE tunnel must also be configured on the OmniAccess WLAN controller. An IP address is required to bring the interface up; if necessary, specify a dummy IP address when configuring the interface. Follow the steps below to configure a switch for mDNS.

1. Click the Create icon **+** to go to the "Add mDNS Configuration" Screen.
2. Click on the **Select Devices** button to bring up a list of switches that are available for mDNS configuration.
3. Select a switch and click **OK**. You will be returned to the Add mDNS Configuration Screen and the switch you selected will appear in the **Switch** field.
4. Complete the fields as described below:


- **Switch** - The IP address of the selected switch is pre-filled in this field.
- **Select OAW Controller** - Select an OmniAccess WLAN Controller to which you want to connect.
- **mDNS GRE Tunnel** - Select a GRE Tunnel to be used to forward packets to the selected OmniAccess WLAN Controller. This is the GRE Tunnel from the OmniSwitch to the OmniAccess Controller. (Only Layer 2 GRE tunnels are supported.)
- **mDNS Admin Status** - The mDNS administrative status (Enabled/Disabled).
- **Router IP Address** - The router IP address.
- **Router IP Mask** - The router IP mask.
- **Tunnel Source IP Address** - The source IP address of the GRE Tunnel from the switch to the controller. This field is pre-filled with the IP address of the selected switch.
- **Tunnel Dest IP Address** - The destination IP address of the GRE Tunnel from the switch to the controller. This field is pre-filled with the IP address of the selected OAW Controller.
- **VRF ID** - The VRF of the GRE Tunnel. *This field will only be visible if the device supports the Multiple VRF feature **and** SNMPv3 was used for discovery.*

5. Click **OK**. The switch will appear in the list of switches on the mDNS Configuration Screen.

Editing an mDNS Configuration

To edit an mDNS configuration on a switch, select the switch in the mDNS Configuration Table and click on the Edit icon  to bring up the "Edit mDNS Configuration" Screen. Edit the allowable fields as described above and click **OK**. The new configuration will be displayed in the mDNS Configuration Table.

Deleting an mDNS Configuration

To delete an mDNS Configuration, select the switch in the mDNS Configuration Table, click on the Delete icon , then click **OK** at the confirmation prompt.

Viewing mDNS Configurations

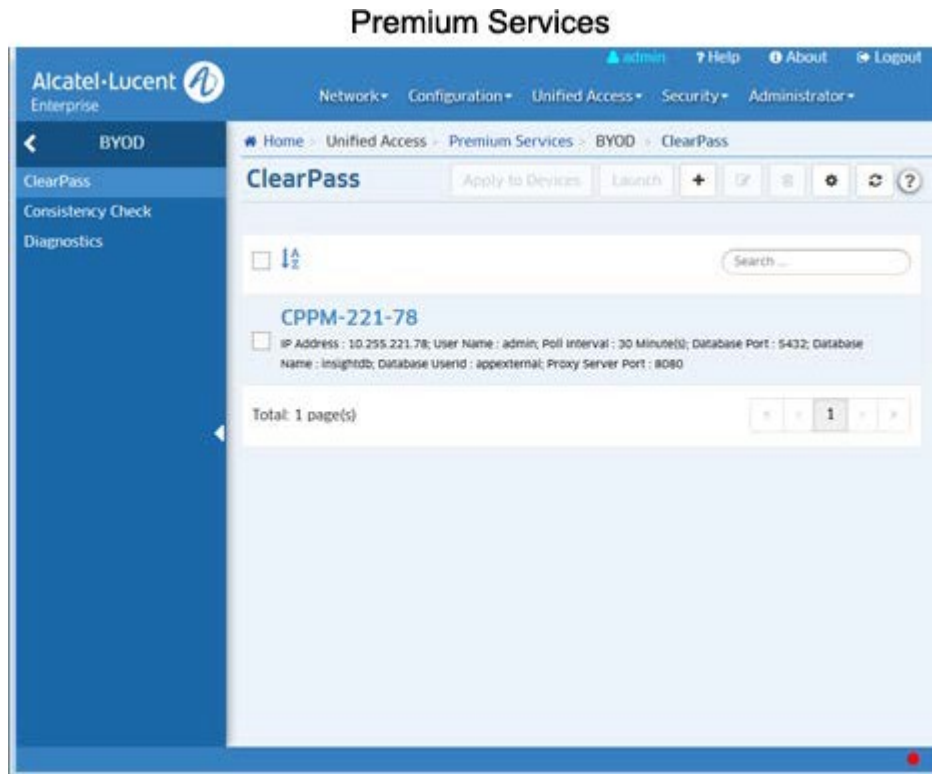
The mDNS Configuration Table displays all switches currently configured for mDNS. The fields are defined below.

- **Name** - The user configured name for the device.
- **Device** - The IP address of the device.
- **mDNS GRE Tunnel** - The name of the GRE Tunnel to be used to forward packets to the OmniAccess WLAN Controller.
- **mDNS Admin Status** - The mDNS administrative status (Enabled/Disabled).
- **mDNS Oper Status** - The mDNS operational status (Up/Down).
- **Router IP Address** - The router IP address.
- **Router IP Mask** - The router IP mask.
- **Admin Status** - The router administrative status (Enabled/Disabled).
- **Tunnel Source IP Address Type** - Currently, only IPv4 is supported.
- **Tunnel Source IP Address** - The source IP address of the GRE Tunnel from the switch to the controller. This field is pre-filled with the IP address of the selected switch.
- **Tunnel Dest IP Address Type** - Currently, only IPv4 is supported.
- **Tunnel Dest IP Address** - The destination IP address of the GRE Tunnel from the switch to the controller. This field is pre-filled with the IP address of the selected OAW Controller.

Premium Services

The Unified Access Premium Services application includes the Alcatel-Lucent Enterprise OmniSwitch implementation of Bring Your Own Device (BYOD). BYOD leverages Access Guardian features along with ClearPass Policy Manager (CPPM) to allow a wired or wireless guest, device, or authenticated user to connect to the network through an OmniSwitch edge device using the CPPM Server for unified authentication.

Note BYOD is only supported on Alcatel-Lucent Enterprise Switches running AOS 6.4.6.R01 and later, AOS 6.6.5.R01 and later, AOS 7.3.4.R02 and later, and AOS 8.1.1.R01 and later. OmniVista supports CPPM Manager 6.2, 6.3, and 6.4. Additional configuration of the CPPM Server per CPPM's documentation is required for the BYOD solution to work.

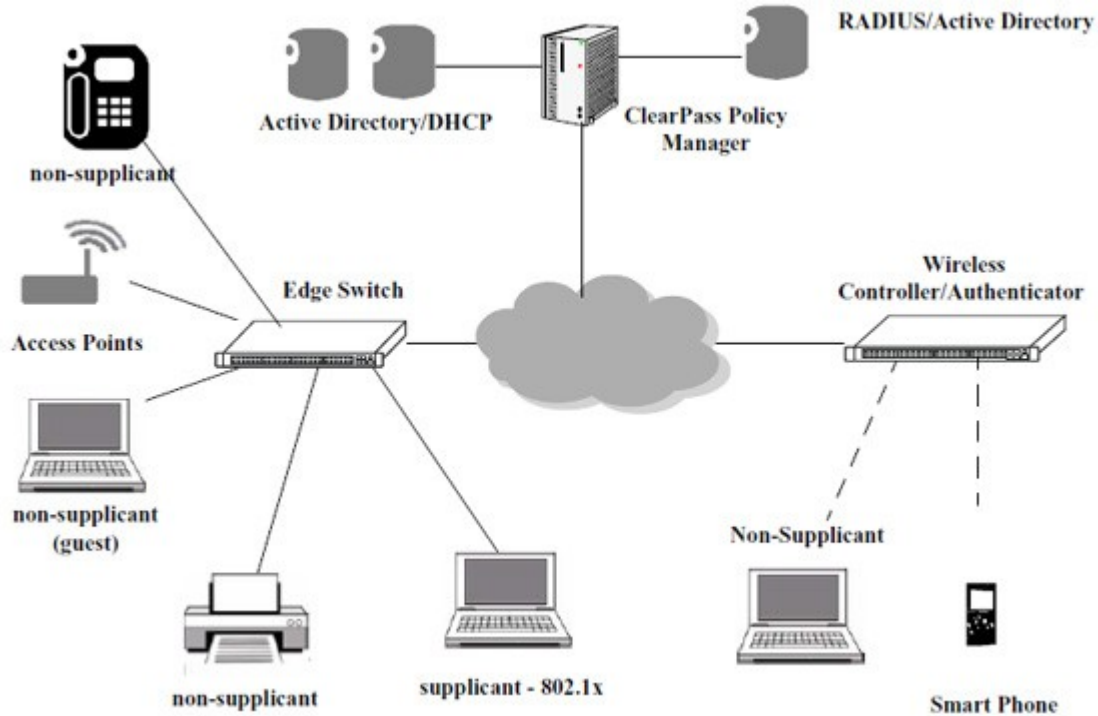


BYOD/ClearPass Overview

The Alcatel-Lucent Enterprise OmniSwitch implementation of BYOD leverages Access Guardian features along with ClearPass Policy Manager (CPPM) to allow a wired or wireless guest, device, or authenticated user to connect to the network through an OmniSwitch edge device using the CPPM Server for unified authentication. The Unified Access application in OmniVista is used to connect to and configure the CPPM Server as a RADIUS Server. Device authentication and network access policies are configured on the ClearPass Server through the ClearPass software and directed to the desired network resources using the UNP feature in the Access Guardian application. The OmniSwitch BYOD solution comprises of the following main components:

- The network infrastructure consisting of both wireless and wireline network. OmniSwitch leverages the Access Guardian features such as 802.1x supplicants, non-supplicant MAC authentication, and User Network Profiles (UNP) to support the BYOD solution.
- The CPPM interacts with both wireless and wireline networks acting as a RADIUS server or RADIUS server proxy. The CPPM provides policy management, guest access, onboarding, and posture checking capabilities.

The figure below provides a high-level view of a typical BYOD Network configuration. The BYOD/ClearPass setup is detailed below.



BYOD Authentication Process Overview

This section describes the basic BYOD process with respect to the OmniSwitch and its interaction with the ClearPass Server.

Authentication for Registered Devices (802.1x)

The BYOD solution provides the following authentication process for registered devices (for example, IT issued employee devices):

1. When an 802.1x enabled port on OmniSwitch detects the user the authentication process is triggered to classify the user.
2. The OmniSwitch sends a request to the ClearPass Server, which authenticates the user based on the user credentials and returns the Role/UNP configured on the ClearPass Server.
3. The ClearPass Server authenticates the user based on the user credentials, and returns the Role/UNP configured on the server to the OmniSwitch.
4. The OmniSwitch assigns the user to the UNP obtained from the ClearPass Server.

Authentication for Network Devices (MAC Authentication)

The BYOD solution provides the following MAC authentication process for network devices such as IP phones, printers, or access points.

1. When MAC authentication is enabled on a port and the OmniSwitch detects the device, MAC authentication process is triggered to classify the device.
2. The OmniSwitch sends a request to the ClearPass Server that authenticates the device based on the devices MAC address and the profiles and policies configured on the ClearPass Server.
3. ClearPass classifies the device to a UNP and returns the UNP information to the OmniSwitch.
4. The OmniSwitch assigns the device to the UNP obtained from the ClearPass Server.

Authentication for Guest Devices and Employees Onboarding

The BYOD solution provides the following authentication process for guest devices and employee personal devices:

1. When MAC authentication is enabled on a port and the OmniSwitch detects the device, the MAC authentication process is triggered to classify the device.
2. ClearPass initially classifies the device to a temporary UNP and returns a redirection URL that allows for guest registration or employee onboarding.
3. OmniSwitch assigns the user to the specified UNP. Since redirection is also set, all DHCP or DNS traffic is allowed but HTTP traffic from the user is redirected towards the URL returned in the UNP.
4. The user is presented with a guest login page or an onboarding page to enter user credentials.
5. ClearPass determines the appropriate role of the user after doing registration and sends the final UNP to the OmniSwitch through a CoA request or RADIUS packet for the case of onboarding.

BYOD/ClearPass Setup

As mentioned earlier, the Alcatel-Lucent Enterprise OmniSwitch implementation of BYOD leverages Access Guardian features on the OmniSwitch along with ClearPass Policy Manager (CPPM) to authenticate users onto the network. The first step in configuring BYOD is to set up the ClearPass Server. The following key points must be considered when configuring the CPPM Server and OmniSwitch for BYOD integration:

- ClearPass Policies and the ClearPass Database must be configured on the CPPM Server using the ClearPass Policy Manager Application.
- The CPPM Server connection, including ClearPass Database login information must be configured in OmniVista so that OmniVista can connect to the CPPM Server and access the CPPM Database.
- The UNP Profile Name configured in OmniVista must be the same as the Enforcement Policy name configured in ClearPass.
- The Insight Database must be enabled on the ClearPass Server for OmniVista to gather Locator information. In the CPPM Application go to: **Administration - Server Configuration**, then click on the Server in the table to bring up the following screen. Make sure the **Enable Profile** and **Enable Insight** checkboxes are checked.

ClearPass Policy Manager

Administration » Server Manager » Server Configuration - alu-cppm
Server Configuration - alu-cppm (10.255.95.250)

System	Services Control	Service Parameters	System Monitoring	Network
Hostname: <input type="text" value="alu-cppm"/> Policy Manager Zone: <input type="text" value="default"/>				
Enable Profile: <input checked="" type="checkbox"/> Enable to allow this server to perform endpoint classification Enable Insight: <input checked="" type="checkbox"/> Enable Insight on this server				
Management Port:				
IP Address:	<input type="text" value="10.255.95.250"/>			Data/External Port:
Subnet Mask:	<input type="text" value="255.255.255.0"/>			
Default Gateway:	<input type="text" value="10.255.95.254"/>			
DNS Settings:				
Primary				
IP Address:	<input type="text" value="198.206.181.70"/>			Secondary
AD Domains: Policy Manager is not part of any domain. Join to domain here. <input type="button" value="Join AD Domain"/>				

Note: ClearPass Policies are configured on the ClearPass Server using the CPPM application web interface. The CPPM web interface can be accessed by entering the CPPM Server IP address into a browser or configuring the CMMP Server connection in OmniVista and clicking the **Launch** button. The procedures below provide steps to configure OmniVista to interface with the CPPM Server to direct users to the proper UNP following ClearPass authentication. Detailed ClearPass Policy configuration instructions are included in the ClearPass online help. An overview of BYOD and sample ClearPass policy configurations are available in Chapter 43 - "Configuring Access Guardian", in the *OmniSwitch Network Configuration Guide*.

Quick Steps to Configure OmniVista for BYOD

Both the Unified Access application and the Access Guardian application (UNP) are used to configure OmniVista for BYOD. You first use the BYOD application to configure the ClearPass Server connection to OmniVista, and to configure the ClearPass Server as a RADIUS Server. You then use the Access Guardian Application to create UNP policies to provide the user with access to the proper network resources. The sections below provide "quick steps" to initially configure OmniVista for BYOD. For more detailed procedures, click here.

Unified Access Configuration

After setting up the ClearPass Server, follow the steps below to configure the OmniVista connection to the CPPM Server (Management and Database Sections) and to configure the ClearPass Server as a RADIUS Server. You can also configure a Redirect URL for Guest user login.

Note: Some key fields on the "Add ClearPass Server" screen are pre-filled with default values. It is recommended that you do not change the pre-filled values.

1. In the Unified Access Application, click on the ClearPass Link, then click on the Create+ icon to bring up the "Add ClearPass Server" screen.
2. In the **Management** fields, enter the **IP Address**, **User Name**, and **Password** of the CPPM Server. The remaining fields are pre-filled with the default values.
3. In the **Database** fields, enter the **Database Password** (the password should match the one configured when the Database Connection was set up in ClearPass). The remaining fields are pre-filled with the default values.
4. In the **Radius Server** fields, enter the **CPPM Server Shared Secret** (with confirmation). You can also complete the **Backup IP Address/Host Name** field, if applicable.
5. For web authentication, complete the **Redirect Options** fields to specify **Proxy Server Port** and add the **Allowed Server** IP address and mask.
6. When you have completed all of the fields, click the **Apply** button. The Server will now appear in the ClearPass list of servers.
7. Select the server and click the **Apply to Devices** button at the top of the screen to assign the server to specific switches.
8. Configure the fields on the Assignment Screen as described below.

- **ClearPass:** The name and IP address of the CPPM Server (informational only, not configurable).
- **Vendor Name:** Select a vendor to automatically enable the RADIUS dictionary associated with that vendor. Note that the IETF dictionary containing the standard the set of RADIUS attributes is always loaded and is generally used. (Default = IETF)
- **CoA Port:** The UDP port used to send CoA actions (Default = 3799).

9. Select the switch(es) to which you want to assign the CPPM Server, then click the **Apply** button.

Note: The "Available Switches" area will be populated with all of the switches available on the network. However, BYOD is only supported on Alcatel-Lucent Enterprise Switches running AOS 6.4.6.R01 and later, AOS 8.1.1.R01 and later.

Access Guardian Configuration

After configuring the ClearPass Server, go to the Access Guardian application to create UNP policies to provide authenticated users access to the proper network resources.

The screenshot displays the 'Create Edge Profile' configuration page in the Alcatel-Lucent NMS-E. The page has a blue header with the Alcatel-Lucent logo and navigation tabs for Network, Configuration, Unified Access, Security, and Administrator. The main content area is titled 'Edge Profile' and contains the following fields and controls:

- Profile Name:** A text input field with an asterisk indicating it is required.
- Auth Flag:** A radio button set with 'DISABLED' selected.
- Mobile Tag Status:** A radio button set with 'DISABLED' selected.
- Redirect Status:** A radio button set with 'DISABLED' selected.
- Policy List:** A dropdown menu with a plus sign to the right.
- Captive Portal Auth:** A radio button set with 'DISABLED' selected.
- Captive Portal Profile:** A dropdown menu with a plus sign to the right.

At the bottom right of the form, there are 'Apply' and 'Cancel' buttons.

1. On the **Edge Profile Screen**, create an Edge (UNP) Profile (e.g., UNP-employee). The Profile Name **must** match the profile name configured in the Enforcement Profile Screen in ClearPass. If required for the policy, select a Policy List from the **Policy List Name** drop-down field.

2. Enable **Redirect Status** to allow users to be redirected by CPPM while in this UNP.

ClearPass

The BYOD ClearPass Screen displays all configured ClearPass Policy Manager (CPPM) Servers and is used to configure the connection to the CPPM Server, configure the CPPM server as a RADIUS Server, and assign the CPPM server to switches on the network. The screen is also used to edit and delete servers. Once the CPPM Server is configured, you can also launch the ClearPass Policy Manager Web interface by selecting a server and clicking on the **Launch** button.

ClearPass provides Bring Your Own Device (BYOD) access to the network. The Alcatel-Lucent Enterprise BYOD solution integrates with ClearPass Policy Manager (CPPM), using the RADIUS (RFC 3576) Change of Authorization (CoA) to achieve this functionality. OmniVista 2500 NMS supports some portions of the configuration to facilitate the solution, including:

- Providing access to end device information from CPPM using Locator.
- Pairing CPPM and switches to enable RADIUS authentication request/response and RADIUS CoA messages between switches and CPPM.
- Setting up authentication servers for 801.1x, MAC, and Accounting Servers on switches to point to the CPPM Server.
- Enabling switches to accept a Redirection URL contained inside the returned RADIUS message so that HTTP/HTTPS traffic can be redirected to a guest registration/onboard portal.
- Enabling switches to accept the Redirection URL not just from the CPPM Server, but from other remediation portals for host integrity compliance on AOS 8.1:1 Switches using the Allowed Server(s) configuration.

Configuring a ClearPass Server

Click on the Create icon **+**. Complete the fields as described below, then click on the **Create** button. When you are finished, select the checkbox next to the server and click on the **Apply to Devices** button to assign the server to switches on the network. Note that certain key fields are pre-filled with default values. It is recommended that you use the default values for these fields.

Management

- **Server Name:** The CPPM Server name (pre-filled with the default **ClearPass Server**).
- **IP Address:** The IP address of the CPPM Server.
- **User Name:** The Administrative login for the CPPM Server.
- **Password:** The password for the CPPM Server.
- **Poll Interval:** The number of Minutes/Hours/Days to poll the CPPM Server for information (pre-filled with default of **30**), and the poll interval unit: Minutes/Hours/Days (pre-selected with default **Minutes**).

Database

- **Database Port:** The port used to connect to the CPPM Database (pre-filled with default **5432**).
- **Database Name:** The CPPM Server Database name (pre-filled with default **insightdb**).
- **Database User ID:** The CPPM Server Database username (pre-filled with default **appexternal**).
- **Database Password:** The CPPM Server Database password.

RADIUS Server

- **Shared Secret:** The CPPM Server Shared Secret.
- **Confirmed Shared Secret:** Re-enter the CPPM Server Shared Secret.
- **Backup IP Address/Host Name:** The CPPM Server can optionally have a backup server. If you wish to define a backup server that will be used if this server is unavailable, enter the IP address OR Host name of the backup server.
- **Timeout:** The number of seconds a switch will wait before a request to the CPPM Server is timed out.
- **Retries:** The number of retries a switch will attempt when trying to contact the CPPM Server.
- **Authentication Port:** The port you to access the server.
- **Accounting Port:** The port used for accounting information.

Redirect Options

The redirect option is only supported on OS6860 Switches (AOS 8.1.1.R01 and later).

- **Proxy Server Port:** The HTTP proxy port number to use for redirection for Guest user login. (Range = 1024–49151, Default = 8080) server.
- **Allowed Server:** The In addition to the CPPM server, other servers can be added to the 'Allowed Server' list to enable additional redirections for Health and Posture checks.

Note: The Insight Database must be enabled on the ClearPass Server for OmniVista to gather Locator information. In the CPPM application, go to: **Administration - Server Configuration**, then click on the Server in the table to bring up the following screen. Make sure the **Enable Profile** and **Enable Insight** checkboxes are checked.

The screenshot displays the 'ClearPass Policy Manager' interface for 'Server Configuration - alu-cppm (10.255.95.250)'. The navigation path is 'Administration > Server Manager > Server Configuration - alu-cppm'. The configuration page is divided into several sections:

- System:** Hostname: alu-cppm
- Services Control:** Policy Manager Zone: default
- Service Parameters:**
 - Enable Profile: Enable to allow this server to perform endpoint classification
 - Enable Insight: Enable Insight on this server
- System Monitoring:** Management Port:
 - IP Address: 10.255.95.250
 - Subnet Mask: 255.255.255.0
 - Default Gateway: 10.255.95.254
- Network:** Data/External Port:
 - DNS Settings: Primary: 198.206.181.70
 - Secondary: (empty)
- AD Domains:** Policy Manager is not part of any domain. Join to domain here. [Join AD Domain](#)

Assigning a ClearPass Server


OmniVista will configure CPPM as a RADIUS Server on the selected switches. It also sets 802.1x authentication, MAC authentication, and accounting to point to the CPPM RADIUS Server entry. It also sets Redirect Server to be the CPPM Server, allowing the switch to accept redirect messages from the RADIUS Server for Captive portal (Web) authentication with CCPM. On CPPM, the selected switches will be added to the list of Network Access Devices (NAD) with the CoA flag and CoA port. The result is the successful

pairing of the CPPM Server and the switches. If specified, Allowed Servers are also configured on AOS 8.1.1 switches.


When you click the **Apply to Devices** button, the Assignment Screen appears. Configure any options, as described below, then use the "Assign Switch" **Add/Remove** buttons to select the switch(es) and click **Apply** or **Override**. ("Override" will override any previous configurations.)

- **ClearPass:** The name and IP address of the CPPM Server (informational only, not configurable).
- **Vendor Name:** Select a vendor to automatically enable the RADIUS dictionary associated with that vendor. Note that the IETF dictionary containing the standard the set of RADIUS attributes is always loaded and is generally used. (Default = IETF)
- **CoA Port:** The UDP port used to send CoA actions (Default = 3799).
- **Enable Endpoint Profiling** - Enables/Disables the IP Helper function on switches assigned to a ClearPass Server so DHCP request information can be forwarded to ClearPass for endpoint profiling.

Editing a ClearPass Server

Select the ClearPass Server you want to edit and click on the Edit icon  to bring up the "Edit ClearPass Server" Screen. Edit the fields as described above then click on the **Save** button to save the changes to the server.

Deleting a ClearPass Server

To delete a ClearPass Server(s), select the server(s) in the table and click on the Delete icon , then click **OK** at the confirmation prompt.

Launching the ClearPass Web Interface

Once the connection to the CPPM Server has been configured, click on the Launch button to launch the ClearPass Policy Manager web interface. This is where you will configure ClearPass authentication and network access policies. See the ClearPass Policy Manager On-Line Help for ClearPass configuration information.



Consistency Check

To validate BYOD configuration, the Consistency Check Screen is used to ensure that switches configured for ClearPass have all UNPs and Policy Lists required to provide authenticated users with access to the proper network resources. When a Consistency Check is performed, OmniVista searches through list of configured switches and enables the user to run a consistency check on each switch. When the check is complete, a summary screen appears listing any missing UNPs or Policy Lists.

Note: For 8.1.1 switches, which support Policy List overwrite, OmniVista will check to see if the switches have all of the same UNPs and Policy lists as the ClearPass Server. For 6.4.6 switches, which do not support Policy List overwrite, OmniVista will only check to see if the switch has the same UNPs.

Performing a Consistency Check

To perform a consistency check, select a ClearPass Server in the **ClearPass IP** drop-down menu, then click on the **Browse** button to display any configured switches for that ClearPass Server. Select a switch from the Device Selection window and click **OK**. Click on the **Check** button to run the Consistency Check. OmniVista will search the selected switch and display any ClearPass UNPs and/or Policy Lists that are missing/unassigned on the switch. You can click on the "Create" link next to a parameter to go to the Edge Profile and/or Unified Policy List configuration pages for create a profile/policy list; or click on the "Assign" link to assign the profile/policy list to a switch.

Authentication Records

The BYOD Authentication Records Screen displays ClearPass Authentication Records. The screen can be accessed from the Locator application by selecting a MAC Address in the Netforwarding Table to view records for that MAC Address. You can also to view records for a specific MAC Address by clicking on the Configuration button at the top of the page and entering the applicable information.

Authentication Record Fields

- **Timestamp** - The date and time the information was gathered.
- **User Name** - The name used to authenticate.
- **Authentication Status** - The status of authentication.
- **Enforcement Profile** - The name of the ClearPass Policy Manager (CPPM) Enforcement Profile.
- **MAC Address** - The MAC address of the device directly connected to the end station.
- **IP Address** - The IP address of the endpoint.
- **Services** - CPPM's components that serve specific types of requests (e.g., 802.1X Wired, MAC Authentication, Web-based Authentication).
- **Role** - The CPPM role assigned to the authenticated user.
- **Authentication Source Used** - The type of authentication user (e.g., Local User, Guest User).
- **NAD IP** - The IP address of the Network Access Device (NAD) sending network access requests to CPPM Policy Manager.
- **NAD Port ID** - The NAD port sending network access requests to CPPM Policy Manager.
- **NAS Port Type** - The Network Access Server Port type used to authenticate subscribers (e.g., async, cable, ethernet).
- **NAS Identifier** - String identifying the device originating the access request to the CPPM Server.
- **SSID** - The wireless network name.
- **VLAN** - The device VLAN.

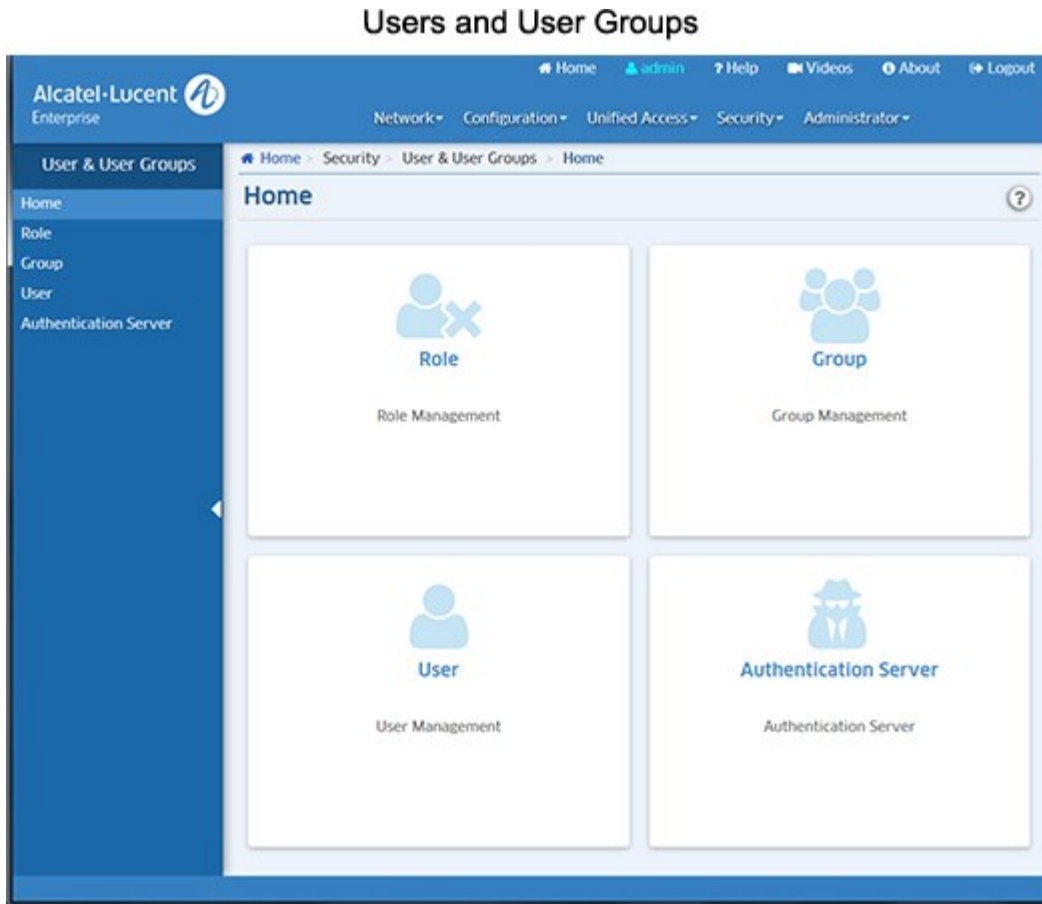
- **Error Code** - The authentication error code.
- **Protocol** - The protocol used to authenticate (e.g., RADIUS or Application) .
- **SPT** - System Posture Token - token returned as the evaluation of health of the client:
 - **Healthy** - Client is compliant: there are no restrictions on network access.
 - **Checkup** - Client is compliant; however, there is an update available. This can be used to proactively remediate to healthy state.
 - **Transient** - Client evaluation is in progress; typically associated with auditing a client. The network access granted is interim.
 - **Quarantine** - Client is out of compliance; restrict network access, so the client only has access to the remediation servers.
 - **Infected** - Client is infected and is a threat to other systems in the network; network access should be denied or severely restricted.
 - **Unknown** - The posture token of the client is unknown.
- **Session ID** - The session identification of authentication.

Authentication Records Configuration

The BYOD Authentication Records Configuration Screen is used to view ClearPass Authentication Records for a device/MAC Address. Enter the IP address of the ClearPass Server in the **Server** Field, select **IP Address** or **MAC Address** in the **Search By** field, enter the address in the **IP Address/MAC Address** field, and click on the **Apply** button. You will be returned to the Authentication Records Table with the information for the selected device displayed.

25.0 Users and User Groups

The Users and User Groups application enables you to control user access to OmniVista and to network devices. Access to OmniVista is controlled through the definition of user logins and passwords. Access to network switches is controlled through the use of User Groups, which have specified levels of access to switches. You can further define access with the User Role feature, which can be used to specify read/write access to specific OmniVista applications and network devices. All OmniVista users must be assigned to at least one User Group, which defines the access rights and roles for its members. User Groups and user logins are configured from the Users and User Groups application, and constitute one level of network security. Other levels of security are summarized below.



User Groups, Users, and User Roles are configured using the following screens:

- **Role Management** - Used to configure User Roles to restrict user access/rights to specific devices and OmniVista applications.
- **Group Management** - Used to configure User Groups to define access to OmniVista, network devices. A User Role is associated with a User Group to specify read/write access to specific devices and OmniVista applications.
- **User Management** - Used to configure users and assign the user to a User Group.
- **Authentication Server** - Used to specify the OmniVista Login Server.

Note: A User Role is an option that enables you to provide user access/rights to specific applications and network devices. For the most part, configuring Users and User Groups is all that will be required.

Security Levels

Security levels are configured in the Users and User Groups application, and through the Command Line Interface (CLI):

- **SNMP Get and Set Community Names** - Get and Set Community names act as read and write passwords that define whether any OmniVista user is allowed to read or write the switch's configuration information. Get and Set Community names are configurable only from the switch itself. Configured through the Console Port or CLI.
- **OmniVista User Groups** - User Groups in OmniVista provide different level of access to switches. An OmniVista user's access rights are based on the access rights of his/her assigned User Group. Configured in the Users and User Groups application.

Default Groups, Users, Roles

OmniVista security uses a combination of user logins, User Groups, and User Roles to control access to OmniVista, network switches, and applications. OmniVista is shipped with the pre-configured user logins, passwords, and User Groups described below. The Users and User Groups application enables you to modify these User Groups, Users, and passwords, or create new ones. Note that the pre-configured user **admin** is the only user that has permission to change the user logins and User Groups defined by the Users and User Groups application. The pre-configured User Groups, Users, and Roles shipped with OmniVista are as follows:

Group	User	Role	Access
Administrators	admin	Account Admin	Full administrative rights to all devices in the network and full administrative rights to the following features. These features are only available to this user: <ul style="list-style-type: none"> • User Management • License Management • Write Operations of System Settings • Control Panel Watchdog, Scheduler Management, and Session Management The default password for this user is switch .
Network Administrators	netadmin	Network Admin	Full administrative rights to all devices in the network. The default password for this user is switch .
Writers	writer	Write	Read/Write access to all devices in the network. The default password for this user is switch .
Default	user	Read	Read access to all devices in the network. The default password for this user is switch .

Note: A User Role is an option that enables you to provide user access/rights to specific OmniVista applications and network devices. For the most part, configuring Users and User Groups is all that will be required. The User Roles feature is configured on the Role Management Screen. This feature enables you to specify access to specific applications, as well as devices using Topology maps. You can also limit user access to specific devices for VLAN and VXLAN configuration. You create a User Role to specify user access, associate it with a User Group, and then create a user in that User Group.

Working with User Groups, Users, and User Roles

You can use one of the pre-configured User Groups or use the Group Management Screen to create a new group or edit one of the pre-configured groups. You can use one of the pre-configured users or use the User Management Screen to create a new user or one of the edit pre-configured users. And you can use the Role Management Screen to create a new role.

Note: All pre-configured users have the same default password, **switch**. At a minimum, it is recommended that you redefine the passwords.

The User Role feature allows you to limit users to specific network devices and applications. For example, OmniVista users with Admin rights can view and manage every device in the network, and have read/write access for all applications. With the User Role feature, you can limit the devices a user can manage and the applications the user can configure by creating a User Role with access to a specific Topology map.

To utilize the User Role feature, you create a User Role with access to a specific Topology map and read/write access to a specific application(s). You then create a User Group and associate that group with that User Role. And finally, you create a user and associate it with that User Group. The user would then have full administrative rights to the specified applications for all devices in the specified map

For example, you could create a User Role (User Role 1) with access to devices in Map 1 and read/write access to the Application Visibility application. A user with this role would be able to access all devices in Map 1 and configure Application Visibility on those devices. And since a user can have multiple roles, you could create a second User Role (User Role 2) with access to Map 2 and read/write access to the CLI Scripting and assign it to the same user. That user could now configure Application Visibility on devices in Map 1, and CLI Scripting on devices in Map 2.

Role Management

The Users and User Groups Role Management Screen displays all currently-configured User Roles. The screen is used to create, edit, or delete User Roles. The User Role feature enables you to specify user rights for specific OmniVista applications and devices. A User Role is associated with a User Group to define access for users assigned to the group. OmniVista is shipped with four pre-configured User Roles:

- **Account Admin** - This User Role can access all maps and has full administrative access rights to all devices in the network. This User Role also has full administrative rights to edit the groups and users defined in the Users and Groups Application.
- **Network Admin** - This User Role can access all maps and has full administrative access rights to all devices in the network. This User Role can only perform "Edit" operations on Topology maps, and does **not** have administrative rights to edit the groups and users defined in the Users and Groups Application.
- **Write** - This User Role can access all maps and has Read/Write access to all devices in the network.
- **Read** - This User Role can access all maps and has Read access to all devices in the network.

Note: Specific rights for each OmniVista application for the above system-defined Roles can be viewed by clicking on a Role in the Existing Roles Table to view the Details window.

Creating a User Role

Click on the Create icon **+** to launch the Role Management Wizard and configure and create a User Role. Complete the fields as described below. Click on the **Next** button to move to the next window. When you are finished, click on the **Create** button.

Role Info and Map Access

Complete the fields below to specify which Topology maps a user can access.

- **Role Name** - Enter a name for the User Role.
- **Description** - Enter an optional description for the User Role.
- **Accessible Maps** - Select an option from the drop-down menu to specify the maps the user can access. The user will only have access to devices in the selected map(s).
- **All Maps** - The user can access all maps.
- **No Maps** - The user cannot access any maps. The user will only have access to non-network OmniVista applications (e.g., Audit, Preferences).
- **Selected Maps** - Select this option, then click on the **Add/Remove Maps** button to select maps the user can access.

Application Access Control

Select the OmniVista application access for the user. Only those applications you configure (either Read or Write access) will be available to the user. By default, Read access is pre-selected for Topology (if map access is configured), System Preferences and Users and User Groups. Read/Write access is pre-selected for User Preferences and Report.


Object Restrictions

Specify the VLANs and or VXLANs the user can access for VLAN/VXLAN configuration. The user will be able to perform VLAN/VXAN operations on these VLANs/VXLANs for devices specified in the Role Info and Map Access window above. This parameter is optional.


Review

Review the configuration. Click on the **Back** button to make any changes.

Editing a User Role

Click on a User in the Existing Users Table and click on the Edit icon . Edit any fields as necessary and/or edit the User Groups at the bottom of the screen to re-assign the User to a different User Group. When you are done, click **Apply**. You will be returned to the User Management Screen. Note that you cannot edit the User Login field. Note that you cannot edit a system-defined User Role.

Deleting a User Role

Select a User(s) in the Existing Users Table, click on the Delete icon , then click **OK**. Note that you cannot delete a system-defined User Role.

Existing Roles Table

The Existing Roles Table displays all configured Users. Click on a User Role in the table for more details.

- **Role Name** - Role Name.
- **Description** - Role Description.
- **System Defined** - Whether the role is a system-defined role or a user-defined role.
- **Accessible Maps** - The maps a user assigned to this role can access.
- **Access Control** - The access/rights to OmniVista applications for a user assigned to this role.

User Role Feature

Basically, the User Role feature allows you to limit users to specific network devices and OmniVista applications. For example, OmniVista users with Admin rights can view and manage every device in the network, and have read/write access for all applications. With the User Role feature, you can limit the devices a user can manage and the applications the user can configure by creating a User Role with access to a specific Topology map and write access to specific applications.

To utilize the User Role feature, you create a User Role with access to a specific Topology map and read/write access to a specific application(s). You then create a User Group and associate that group with that User Role. And finally, you create a user and associate it with that User Group. The user would then have full administrative rights to the specified applications for all devices in the specified map.

For example, you could create a User Role (User Role 1) with access to devices in Map 1 and read/write access to the Application Visibility application. A user with this role would be able to access all devices in Map 1 and configure Application Visibility on those devices. And since a user can have multiple roles, you could create a second User Role (User Role 2) with access to Map 2 and read/write access to the CLI Scripting and assign it to the same user. That user could now configure Application Visibility on devices in Map 1, and CLI Scripting on devices in Map 2.

The table below provides some use case samples for assigning multiple User Roles to a User.

Scenario	User Role 1	User Role 2	User Role 3	Device/Application Access
Using Topology Maps to limit access to devices	Map 1 Read Access for Topology	Map 2 Write Access for Topology	Map 3 Read Access for Topology	Read Access for devices in Maps 1 and 3. Write Access for devices in Map 2
Using a combination of Topology Maps and an application, such as Application Visibility.	Map 1 Read Access for Application Visibility	Map 2 Read Access for Application Visibility	Map 3 Write Access for Application Visibility	Read Access for Application Visibility for devices in Maps 1 and 2. Write Access for Application Visibility for devices in Map 3.
Using a combination of Topology Maps and an Object (VLAN)	Map 1 VLAN 10 Read Access for Application Visibility	Map 2 VLAN 20 Read Access for Application Visibility	Map 3 VLAN 30 Write Access for Application Visibility	Read Access for Application Visibility for devices in Maps 1 and 2; and VLAN configuration allowed on those devices in VLANs 10 and 20. Write Access for Application Visibility for devices in Maps 1 and 2; and VLAN configuration allowed on those devices in VLAN 30.


Group Management

The Users and User Groups Group Management Screen displays all currently-configured User Groups (along with a brief description). You can click on a User Group in the list for more information about the group. The screen is used to create, edit, or delete User Groups; and add or delete Users from a User Group. OmniVista is shipped with four pre-configured User Groups:

- **Administrators** - This User Group has full administrative access rights to all devices in the network AND full administrative rights to edit the groups and users defined in the Users and Groups Application.
- **Network Administrators** - This User Group has full administrative access rights to all devices in the network. Members of this group are the users who are responsible for management of parts of the network (Site Administrators). This group can manually add, delete, or modify devices.
- **Writers** - This User Group has Read/Write access to all devices in the network.
- **Default** - This User Group has Read access to all devices in the network.

Note: Specific rights for each OmniVista Application for the above system-defined Groups can be viewed by clicking on a Group in the Existing Groups Table to view the Details window.


Creating a User Group

Click on the Create icon  and complete the fields as described below. When you are finished, click on the **Create** button.


- **Name** - Enter a name for the group.
- **Description** - Enter an optional description for the group.
- **Assigned Roles** - Select a User Role for the group.
- **User Members** - Select a User(s) for the group.

Note that users may belong to more than one group at a time, in which case their access rights are defined by the most privileged group to which they belong. Also note that you do not have to add users to the User Group at this time. When you create a user, you can add them to any existing User Group as a member. You can also edit a User Group later to add members.

Editing a User Group

Click on a Group on the Group Management Screen to bring up the User Group Detail Screen. Click on the Edit icon . You can edit the **Description**, **Assigned Roles**, and **User Members** fields. When you are done, click the **Apply** button. You will be returned to the Group Management Screen. Note that you cannot edit the Group Name field. Also note that you can only edit the Description field of the Administrators Group.

Deleting a User Group

Select a User Group(s) on the Group Management Screen by clicking in the checkbox, click on the Delete icon , then click **OK**. Note that you cannot delete the Administrators Group or the Default Group.

Existing Groups Table

The Existing Groups Table displays all configured User Groups. Click on a group in the table for more details.

- **Name** - Group Name.
- **Description** - Group Description.
- **Assigned Roles** - The User Roles assigned to the group.


User Management

The Users and User Groups User Management Screen displays all currently-configured Users by login name (along with a brief description). The screen is used to create, edit, or delete Users. Note that a User's access rights are determined by the User Group in which the user is a member. OmniVista is shipped with four pre-configured Users and four pre-configured User Groups. The default password for all four pre-configured Users is **switch**. For security reasons, it is recommended that you redefine the default passwords. The default Users and their default pre-configured User Group memberships are as follows:


- **admin** - This user belongs to the Administrators User Group and has full administrative rights to all switches on the network AND full administrative rights to the Users and Groups Application. The default password for this user is **switch**.
- **netadmin** - This user belongs to the Network Administrators User Group and has full administrative rights to all devices in the network. The default password for this user is **switch**.
- **writer** - This user belongs to the Writers User Group and has Read/Write access to all devices in the network. The default password for this user is **switch**.
- **user** - This user belongs to the Default User Group and has read access to all devices in the network. The default password for this user is **switch**.

Note: Specific rights for each OmniVista Application for the above system-defined Users can be viewed by clicking on a User in the Existing Users Table to view the Details window.


Creating a User

Click on the Create icon  and complete the fields as described below. The **User Login** and **Password** fields are **mandatory**. The user will use this User Login and Password to log into OmniVista. You can complete additional fields (e.g. Name, Description) to provide a more detailed description of the User. Select one of the User Groups at the bottom of the screen to assign the User to a specific User Group and click on the **Create** button. Note that users may belong to more than one User Group at a time, in which case their access rights are defined by the most privileged group to which they belong.

Editing a User

Click on a User in the Existing Users Table and click on the Edit icon . Edit any fields as necessary and/or edit the User Groups at the bottom of the screen to re-assign the User to a different User Group. When you are done, click **Apply**. You will be returned to the User Management Screen. Note that you cannot edit the User Login field.

Deleting a User


Select a User(s) in the Existing Users Table, click on the Delete icon , then click **OK**. Note that you cannot delete the user admin.

Existing Users Table

The Existing Users Table displays all configured Users. Click on a user in the table for more details.

- **Login** - User login.
- **Description** - User Description.
- **Assigned Roles** - The User Roles assigned to the user. This tab also displays the Application Access Control configuration for the user (the user's read/write access for OmniVista applications).
- **Assigned Groups** - The User Group(s) to which the user belongs.

Authentication Server

The Users and User Groups Authentication Server Screen is used to select the Login Authentication Server. You can select the OmniVista Server (Local) or a remote RADIUS Server. If necessary, click on the Add icon  to go the RADIUS Server Management Screen and configure a remote RADIUS Server. (After creating the server you will automatically be returned to the Authentication Server Screen.)

Select the server from the **Authentication Server** drop-down list. If you select a server other than "Local", select an **Authentication Method** for communication between the remote server and OmniVista (PAP, CHAP, MSCHAP, MSCHAPV2, or EAP - MSCHAPV2).

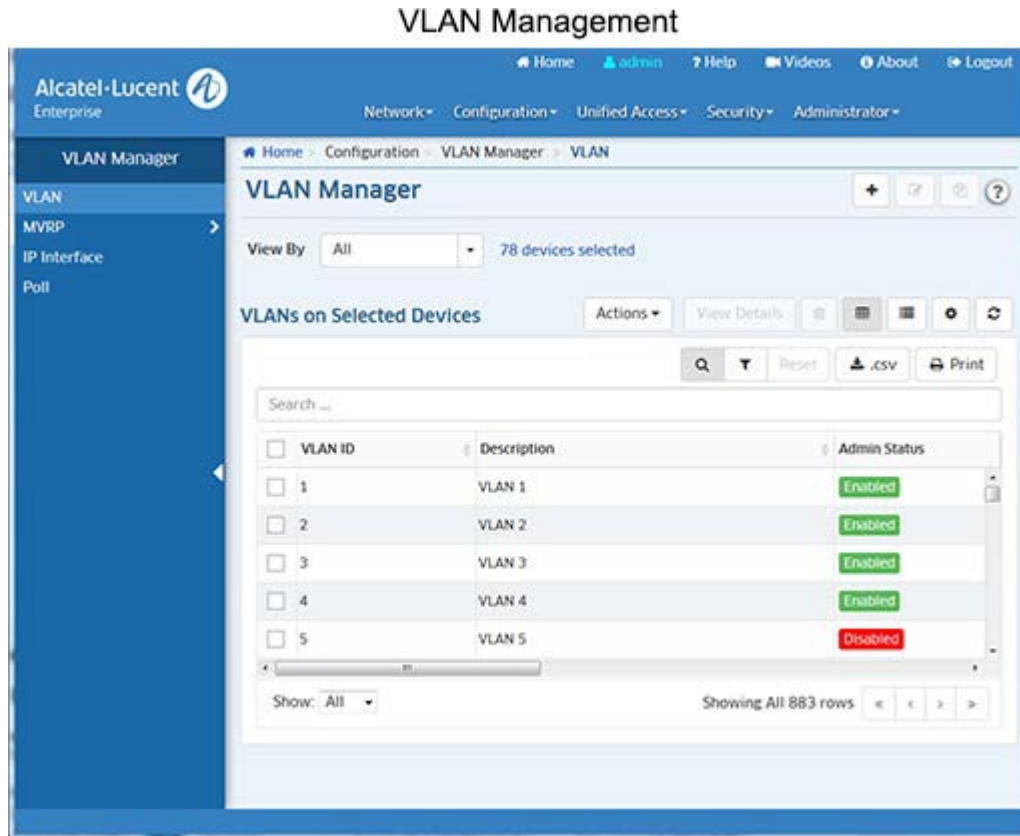
If you select a remote server you must also verify connectivity to that server by clicking on the **Test Radius Server Connection** button and entering your RADIUS User Name and Password for the selected server. (You can also enter the User Name and Password of any user configured on the selected Server.) This is to ensure that the selected RADIUS Server is reachable and is configured correctly for OmniVista remote authentication. Enter your RADIUS Server **User Name** and **Password** and click **OK**. OmniVista will ping the server and verify connectivity and configuration. If successful, the **Apply** button will activate. Click on the **Apply** button to set the new server. If the server is unreachable or not configured correctly, the **Apply** button will not activate and you will not be able to change the server.

Notes:

- Only a remote RADIUS Server or the OmniVista Server (Local Database) can be used for OmniVista login. If a remote Authentication Server is used, and that remote server and the remote backup server are not available, users cannot log into OmniVista. If necessary, the Administrator may change the Authentication Server using the Virtual Appliance Menu on the VA.
- You can also view/change the Authentication Server using the Virtual Appliance Menu on the VA. Open a Console on the VA. On **The Virtual Appliance Menu**, enter **7** to go to the **Login Authentication Server Screen**.
- If a remote authentication server is selected, and that remote server and the remote backup server are not available, users can login from the local OmniVista Server.
- If the Administrator changes the Login Server, current users will remain logged in. However, if the users attempt to login/re-login, they will be logged in using the new Login Server.

26.0 VLAN Manager

The VLAN Manager Screen displays a list of all configured network VLANs; and is used to create, edit, copy, and delete VLANs. It is also used to perform certain actions on a VLAN (e.g., enable/disable VLANs, view/edit Spanning Tree parameters, view/configure an IP Router, and view VLAN details). Links on the left side of the screen are used to view/configure MVRP and IP interfaces; and poll network devices.



Note: The VLAN Manager application is supported on AOS Switches and OAW Controllers.

VLAN Overview

One of the main benefits of using VLANs to segment network traffic is that VLAN configuration and port assignment is handled through 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 initial configuration for all Alcatel-Lucent Enterprise (ALE) switches consists of a default VLAN 1 and all device ports are initially assigned to this VLAN. If additional VLANs are not configured on the switch, then the entire switch is treated as one large broadcast domain. All ports will receive all traffic from all other ports.

In compliance with the IEEE 802.1Q standard, each VLAN is identified by a unique number, referred to as the VLAN ID. The user specifies a VLAN ID to create, modify or remove a VLAN and to assign switch ports to a VLAN. When a packet is received on a port, the VLAN ID for that port is inserted into the packet. The packet is then bridged to other ports that are assigned to the same VLAN ID. In essence, the VLAN broadcast domain is defined by a collection of ports and packets assigned to its VLAN ID.

The operational status of a VLAN remains inactive until at least one active switch port is assigned to the VLAN. This means that VLAN properties, such as Spanning Tree and/or router interfaces, also remain

inactive. Ports are considered active if they are connected to an active network device. Non-active port assignments are allowed, but do not change the VLAN's operational state.

When using the OmniVista VLAN Manager application to configure VLANs in your network, consider the following:

- There is no staging of VLAN configuration changes. When you make a change to a VLAN, changes are sent directly to the device and are processed in real time.
- If an error occurs when changes are applied to a device, any changes successfully made to that point are maintained and not backed out of the switch configuration.
- The parameter values displayed in the VLANs Table, except for the VLAN ID field, are the values obtained from the switch polled that has the lowest IP host address. For example, if VLAN 5 exists on three different switches with IP addresses of 10.0.0.1, 10.0.0.2, and 10.0.0.3 and each instance of the VLAN has a different description, the VLAN 5 description from switch 10.0.0.1 is displayed in this window.
- When you modify VLAN parameters, however, the changes are applied across all switches in the topology that have this VLAN configured. For example, if you selected VLAN 10 and changed the description to "Marketing Department", all switches that contain VLAN 10 would receive this new description value.

VLANs Table


The VLANs Table displays all VLANs configured on the network. By default, VLANs configured on all devices are displayed. However, you can filter the view by selecting "Devices" from the **View by** drop-down menu and selecting specific devices. Only VLANs configured on the selected device(s) will be displayed. The VLANs Table displays basic information about each VLAN as shown below. To view detailed information about a VLAN, double click on the VLAN in the VLANs Table or select the VLAN and click the **View Details** button at the top of the VLANs Table.

- **VLAN ID** - In compliance with the IEEE 802.1Q standard, each VLAN is identified by a unique number, referred to as the VLAN ID. This number is assigned by the user at the time the VLAN is created and is not a modifiable parameter. When a network device packet is received on a port, the port's VLAN ID is inserted into the packet. The packet is then bridged to other ports that are assigned to the same VLAN ID. In essence, the VLAN broadcast domain is defined by a collection of ports and packets assigned to its VLAN ID. (Range = 1 - 4094)
- **Description** - A text string up to 32 characters. This parameter defaults to the VLAN ID number (e.g., VLAN 10) if a description is not specified at the time the VLAN is created.
- **Admin Status** - The administrative status of the VLAN (Enabled/Disabled). By default, the administrative status is enabled when a VLAN is created. When a VLAN is administratively disabled, static port and dynamic mobile port assignments are retained but traffic on these ports is not forwarded. However, VLAN rules remain active and continue to classify mobile port traffic for VLAN membership.
- **Oper Status** - The VLAN operational status (Active/Inactive). This parameter is not modifiable; switch software determines if the VLAN is operationally active or inactive and sets the appropriate field value. A VLAN's operational status remains inactive until at least one active switch port is assigned to the VLAN and the VLAN's administrative status is enabled. This means that VLAN properties, such as Spanning Tree or router ports, also remain inactive. Ports are considered active if they are connected to an active network device. Non-active port assignments are allowed, but do not change the VLAN's operational state.
- **VLAN Type** - The type of VLAN is determined at the time the VLAN is created (e.g., Standard, BVLAN, Control BVLAN).

- **Spanning Tree Status** - The Spanning Tree Status (Enabled/Disabled) for the VLAN. When a VLAN is created, an 802.1D standard Spanning Tree Algorithm and Protocol (STP) instance is enabled for the VLAN by default. STP evaluates VLAN port connections to determine if there are redundant data paths between the same VLAN on other switches. If a redundant path does exist, STP determines which path to block in order to provide a loop-free network topology. In this manner, STP ensures that there is always only one active data path between any two switches (VLANs). When a change occurs, such as a path is disconnected or a path cost change, the Spanning Tree Algorithm activates the blocked path to restore the network connection.
- **Router Protocol** - The protocol for the VLAN virtual router port. If no router port is configured for the VLAN, then "none" appears in this field. A VLAN is available for routing when a virtual router port is defined for that VLAN and at least one active port has joined the VLAN. If a VLAN does not have a router port, its ports are in essence firewalled from other VLANs.

Note: The basic information displayed in the VLANs Table, except for the VLAN ID field, is the information obtained from the switch polled that has the lowest IP host address. For example, if VLAN 9 exists on three different switches with IP addresses of 10.0.0.1, 10.0.0.2, and 10.0.0.3 and each instance of the VLAN has a different description, the VLAN 9 description from switch 10.0.0.1 is displayed in this window. You can view detailed information on each device in the VLAN by double-clicking on a VLAN in the VLANs Table, or selecting a VLAN and clicking on the **View Details** button at the top of the VLANs Table.


Creating a VLAN

VLANs are created using a wizard that guides you through each of the steps needed to create the VLAN. To create a VLAN(s), click on the Create icon . The Create VLAN Wizard will then guide you through the following screens:


- Device Selection - Basic VLAN configuration parameters (e.g. VLAN ID, Description, administrative status) and device selection.
- VLAN Configuration - Review VLAN device selection and review/modify VLAN administrative status.
- Default Ports Assignment - Configure VLAN Ports on selected device(s).
- Q-Tagged Ports Assignment - Configure Q-Tagged Ports on selected device(s).
- Review - Review VLAN configuration and create VLAN.

Note: When creating a VLAN, you can select up to 200 devices. If necessary, once the VLAN is created you can edit the VLAN to add additional devices. Again, for each edit, you can add up to 200 devices. Repeat to add additional devices. Also note that you can use the wizard to create a single VLAN or multiple VLANs.

Editing a VLAN



To edit a VLAN, select the VLAN in the VLANs Table and click on the Edit icon . The Edit VLAN Wizard appears. Use the wizard to make any edits. When you are done, go the Review Window in the wizard and click on the **Apply** button to update the VLAN. Note that you cannot edit the VLAN ID.

Copying a VLAN

You can save time creating a new VLAN by copying an existing VLAN and modifying the configuration. Select a VLAN in the VLANs Table and click on the Copy icon . The Create VLAN Wizard will appear. Enter a new VLAN ID and use the wizard to create the new VLAN.

Deleting a VLAN

You can delete a VLAN(s) completely - delete the VLAN(s) and delete them from all devices; or you can delete a VLAN(s) from specific devices.

- To delete a VLAN(s) completely, make sure the **View by** drop-down menu is set to "All". Select the VLAN(s) in the VLANs Table, click on the Delete icon , then click **OK** at the confirmation prompt. The VLAN(s) will be removed from OmniVista and from all network devices to which it was assigned.
- To delete a VLAN(s) from specific devices, select "Devices" from the **View by** drop-down menu and select the device(s). Select the VLAN(s) in the VLANs Table, click on the Delete icon , then click **OK** at the confirmation prompt. The VLAN(s) will be removed from the selected devices. Note that you can also delete a VLAN from a specific device by editing the VLAN and removing the device(s) from the VLAN.

Note: You cannot delete a VLAN from more than 200 devices at a time.

VLAN Actions

You can perform the following actions on VLANs by selecting a VLAN in the VLANs Table and clicking on the **Actions** button at the top of the screen: enable/disable VLANs, view/modify Spanning Tree parameters, view/configure an IP Router, and view VLAN details.


Enabling/Disabling VLANs

To enable/disable a VLAN(s), select one or more VLANs in the table, click on the **Actions** button and select **Enable** or **Disable**. The Results Screen displays the operation results. Click **OK** to return to the VLAN Manager Screen.

Viewing/Modifying Spanning Tree

The Spanning Tree Screen in the VLANs application is used to view Spanning Tree configuration information for devices in a VLAN, and to modify Spanning Tree Bridge or Port Parameters. To view/edit Spanning Tree information for a VLAN, select the VLAN in the VLAN's Table, then select **Spanning Tree** from the **Actions** drop-down menu at the top of the screen.

Viewing/Configuring IP Routers

The IP Router Screen in the VLANs application is used to view/configure IP Routers on a VLAN. To configure an IP router for a VLAN, select the VLAN in the VLAN's Table, then select **IP Router** from the **Actions** drop-down menu at the top of the screen. The IP Router Screen will appear. Click on the  Create icon to create an IP router for the VLAN.

VLAN Details

The VLAN Manager Screen displays all VLANs configured on the network. By default, VLANs configured on all devices are displayed. However, you can filter the view by selecting "Devices" from the **View by** drop-down menu and selecting specific devices. Only VLANs configured on the selected device will be displayed. The VLANs Table displays basic information about each VLAN. To view detailed information about a VLAN, select the VLAN, click on the **View Details** button at the top of the VLANs Table.

Basic Information

The VLANs Table displays basic information about each VLAN.

- **VLAN ID** - In compliance with the IEEE 802.1Q standard, each VLAN is identified by a unique number, referred to as the VLAN ID. This number is assigned by the user at the time the VLAN is created and is not a modifiable parameter. When a network device packet is received on a port, the port's VLAN ID is inserted into the packet. The packet is then bridged to other ports that are assigned to the same VLAN ID. In essence, the VLAN broadcast domain is defined by a collection of ports and packets assigned to its VLAN ID. (Range = 1 - 4094)
- **Description** - A text string up to 32 characters. This parameter defaults to the VLAN ID number (e.g., VLAN 10) if a description is not specified at the time the VLAN is created.
- **Admin Status** - The administrative status of the VLAN (Enabled/Disabled). By default, the administrative status is enabled when a VLAN is created. When a VLAN is administratively disabled, static port and dynamic mobile port assignments are retained but traffic on these ports is not forwarded. However, VLAN rules remain active and continue to classify mobile port traffic for VLAN membership.
- **Oper Status** - The VLAN operational status (Active/Inactive). This parameter is not modifiable; switch software determines if the VLAN is operationally active or inactive and sets the appropriate field value. A VLAN's operational status remains inactive until at least one active switch port is assigned to the VLAN and the VLAN's administrative status is enabled. This means that VLAN properties, such as Spanning Tree or router ports, also remain inactive. Ports are considered active if they are connected to an active network device. Non-active port assignments are allowed, but do not change the VLAN's operational state.
- **VLAN Type** - The type of VLAN is determined at the time the VLAN is created (e.g., Standard, BVLAN, Control BVLAN).
- **Spanning Tree Status** - The Spanning Tree Status (Enabled/Disabled) for the VLAN. When a VLAN is created, an 802.1D standard Spanning Tree Algorithm and Protocol (STP) instance is enabled for the VLAN by default. STP evaluates VLAN port connections to determine if there are redundant data paths between the same VLAN on other switches. If a redundant path does exist, STP determines which path to block in order to provide a loop-free network topology. In this manner, STP ensures that there is always only one active data path between any two switches (VLANs). When a change occurs, such as a path is disconnected or a path cost change, the Spanning Tree Algorithm activates the blocked path to restore the network connection.
- **Router Protocol** - The protocol for the VLAN virtual router port. If no router port is configured for the VLAN, then "none" appears in this field. A VLAN is available for routing when a virtual router port is defined for that VLAN and at least one active port has joined the VLAN. If a VLAN does not have a router port, its ports are in essence firewalled from other VLANs.

Detailed Information

The Detail Screens provide detailed information about devices in the selected VLAN. Click on a link at the top of the screen to display Device, Port, or Link Aggregate information.

Device View

- **Friendly Name** - The IP address of the device.
- **Device Name** - The user-configured name for the device.
- **Device IP Address** - The IP address of the device.
- **Device MAC Address** - The MAC address of the device.
- **Device Version** - The version number of the device software. OmniVista may not be able to determine the software version on some third-party devices. In these cases, the field will be blank.

- **Device Location** - The physical location of the device (user-defined, up to 255 characters). If the user did not specify a location the field displays "Unknown".
- **Device Status** - This field displays the operational status of the
 - device: Up - Device is up and responding to polls.
 - Down - Device is down and not responding to polls.
 - Warning - Device has sent at least one warning or critical trap
- **Device Type** - The device model (e.g., OS6900-X72).
- **VLAN ID** - The VLAN ID number (e.g., VLAN 10).
- **VLAN Description** - A text string up to 32 characters. This parameter defaults to the VLAN ID number (e.g., VLAN 10) if a description is not specified at the time the VLAN is created.
- **VLAN Admin Status** - The administrative status of the VLAN (Enabled/Disabled). By default, the administrative status is enabled when a VLAN is created. When a VLAN is administratively disabled, static port and dynamic mobile port assignments are retained but traffic on these ports is not forwarded. However, VLAN rules remain active and continue to classify mobile port traffic for VLAN membership.
- **VLAN Type** - The type of VLAN is determined at the time the VLAN is created (e.g., Standard, BVLAN, Control BVLAN).
- **Spanning Tree Status** - The Spanning Tree Status for the VLAN (Enabled/Disabled). When a VLAN is created, an 802.1D standard Spanning Tree Algorithm and Protocol (STP) instance is enabled for the VLAN by default. STP evaluates VLAN port connections to determine if there are redundant data paths between the same VLAN on other switches. If a redundant path does exist, STP determines which path to block in order to provide a loop-free network topology. In this manner, STP ensures that there is always only one active data path between any two switches (VLANs). When a change occurs, such as a path is disconnected or a path cost change, the Spanning Tree Algorithm activates the blocked path to restore the network connection.
- **Mobility** - The mobile status for the VLAN (Enabled/Disabled). On AOS switches, mobility is not enabled or disabled at the VLAN level. Instead, switch ports are designated as mobile or non-mobile.
- **Oper Status** - The VLAN operational status (Active/Inactive). This parameter is not modifiable; switch software determines if the VLAN is operationally active or inactive and sets the appropriate field value. A VLAN's operational status remains inactive until at least one active switch port is assigned to the VLAN and the VLAN's administrative status is enabled. This means that VLAN properties, such as Spanning Tree or router ports, also remain inactive. Ports are considered active if they are connected to an active network device. Non-active port assignments are allowed, but do not change the VLAN's operational state.
- **Authentication** - The authentication status for the VLAN (Enabled/Disabled). By default, authentication is disabled when a VLAN is created. Once authentication is enabled on a VLAN, however, then only authenticated mobile port devices can join the VLAN after completing the appropriate log-in process.
Layer 2 authentication uses VLAN membership to grant access to network resources. Authenticated VLANs control membership through a log-in process; this is sometimes called user authentication. A VLAN must have authentication enabled before it can participate in the Layer 2 authentication process.
- **Voice Status** - Administrative status (Enabled/Disabled) of voice usage for the current VLAN (only supported on 6.x devices).
- **Router Protocol** - The protocol for the VLAN virtual router port. If no router port is configured for the VLAN, then "none" appears in this field. A VLAN is available for routing when a virtual router port is defined for that VLAN and at least one active port has joined the VLAN. If a VLAN does not have a router port, its ports are in essence firewalled from other VLANs.

Port View

- **Device Friendly Name** - The IP address of the device.
- **Device IP Address** - The IP address of the device.
- **Port** - The VLAN slot/port number.
- **Port Type** - The Port Type parameter indicates how the port assignment to the current VLAN was made.
 - **Default** - The port is a fixed port that was statically assigned to the VLAN, which is now the configured default VLAN for the port.
 - **Q Tagged** - The port is a fixed port that was statically assigned to the VLAN using the 802.1Q tagging feature. The VLAN is a static secondary VLAN assignment for the 802.1Q tagged port.
 - **Mobile** - The port is a mobile port that was dynamically assigned to the VLAN when traffic received on the port match traffic rules defined for the VLAN. The VLAN is a dynamic secondary VLAN assignment for the mobile port.
- **Port State** - The the status of the VLAN port assignment.
 - **Forwarding** - Port is active and forwarding traffic on this VLAN.
 - **Inactive** - Port is not active (administratively disabled, down, or nothing is connected to the port).
 - **Blocking** - Port is active, but not forwarding any traffic on this VLAN.
 - **Filtering** - Mobile port traffic is filtered for the VLAN; only traffic received on the port that matches VLAN rules is forwarded. Occurs when a mobile port's VLAN is administratively disabled or the port's default VLAN status is disabled. Does not apply to fixed ports.

Link Aggregate View

- **Friendly Name** - The IP address of the device.
- **Device IP Address** - The IP address of the device.
- **Link Aggregate ID** - The ID of the link aggregate group of ports. This number was assigned when the aggregate was created.
- **Link Aggregate** - An optional textual description (up to 32 characters) for the link aggregate.
- **Port Type** - The Port Type parameter indicates how the port assignment to the current VLAN was made.
 - **Default** - The port is a fixed port that was statically assigned to the VLAN, which is now the configured default VLAN for the port.
 - **Q Tagged** - The port is a fixed port that was statically assigned to the VLAN using the 802.1Q tagging feature. The VLAN is a static secondary VLAN assignment for the 802.1Q tagged port.
 - **Mobile** - The port is a mobile port that was dynamically assigned to the VLAN when traffic received on the port match traffic rules defined for the VLAN. The VLAN is a dynamic secondary VLAN assignment for the mobile port.
- **Description** - The standard MIB name for this dynamic aggregate group.

MVRP

Multiple VLAN Registration Protocol (MVRP) provides a mechanism for dynamic maintenance of the contents of dynamic VLAN registration entries for each VLAN, and for propagating the information they contain to other bridges. This information allows MVRP-aware devices to dynamically establish and update their knowledge of the set of VLANs that currently have active members, and through which ports those members can be reached. The main purpose of MVRP is to allow switches to automatically discover some of the VLAN information that would otherwise have to be manually configured.

To configure MVRP, click on the link on the left side of the screen. The MVRP Configuration Wizard guides you through the steps to configure MVRP on network switches/ports.

IP Interface

Network device traffic is bridged (switched) at the Layer 2 level between ports that are assigned to the same VLAN. However, if a device needs to communicate with another device that belongs to a different VLAN, you must configure an IP interface on a device in the VLAN to enable Layer 3 routing to transmit traffic between VLANs. A VLAN is available for routing when at least one IP interface is defined for that VLAN and at least one active port is associated with the VLAN. If a VLAN does not have an IP interface, the ports associated with that VLAN are in essence firewalled from other VLANs.

The IP Interface Screen in the VLANs application is used to view and configure IP interfaces on devices.

Poll

Network devices are polled every 30 minutes for VLAN information update. The VLAN Manager Poll Screen can be used at any time to manually poll devices. Once the selected devices are polled, VLAN information displayed in the VLAN Manager application is be updated for the selected devices.

Spanning Tree

The VLAN Manager Spanning Tree Screen is used to view Spanning Tree configuration information for devices in a VLAN; and to edit Spanning Tree Bridge and Port parameters. By default, a Summary view is displayed. Click on the Bridge or Port link at the top of the to view bridge or port information.

Summary View

The Summary View displays a summary of the Spanning Tree information for devices in the selected VLAN.

- **Device Friendly Name** - The user-defined name for the device.
- **Device IP** - The IP host address that identifies the switch within the management IP network. It is not the same IP host address defined as a virtual IP router port for the selected VLAN, unless the VLAN is part of the management network. In addition, an IP host address appears in this field even if an IP virtual router port does not exist for the VLAN on that particular switch.
- **Protocol** - The VLAN spanning tree algorithm protocol. The algorithm determines the state and role of a port within the spanning tree topology:
 - **STP (802.1D)** - Standard Spanning Tree Algorithm and Protocol (Default).
 - **RSTP (802.1W)** - Rapid Spanning Tree Algorithm and Protocol. RSTP is based on the 802.1D standard algorithm and is designed to provide quick recovery in the event of a link, port or device failure.
 - **MSTP (802.1S)** - Multiple Spanning Tree Protocol. MSTP is an enhancement to 802.1Q Common Spanning Tree Instance (CST). When the switch is running in Flat Mode, a single Spanning Tree instance is applied across all VLAN port connections. MSTP allows the configuration of Multiple Spanning Tree Instances (MSTI) 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. Note that MSTP in VLAN spanning tree view is only displayed for Instance 0 under VLAN 1. None of the other instances will be displayed.
- **Priority** - The bridge priority value (0 - 65535) for the VLAN. The lower the number, the higher the priority value. The bridge priority value is used by the spanning tree algorithm to determine which

VLAN should serve as the root of the spanning tree. (Default = 32768)

- **Maximum Age** - The amount of time, in seconds, that spanning tree information learned from BPDUs received on VLAN ports is retained. When this information has aged beyond the maximum age value, the information is discarded. (Range = (6- 40, Default = 20)
- **Path Cost** - The cost of the path to the root for this Spanning Tree instance.
- **Mode** - The Spanning Tree operating mode for the switch:
 - **Flat Mode (Single Spanning Tree)** - The Spanning Tree Algorithm is applied across all VLANs. For example, if a port belonging to VLAN 10 and a port belonging to VLAN 20 both connect to the same switch, then Spanning Tree Algorithm will block one of these ports.
 - **1x1 (One Spanning Tree per VLAN)** - A single Spanning Tree instance is enabled for each VLAN configured on the switch. For example, if there are five VLANs configured on the switch, then there are five separate Spanning Tree instances. In essence, a VLAN is a virtual bridge in that it will have its own bridge ID and configurable STP parameters, such as protocol, priority, hello time, max age and forward delay. Note: By default, the Spanning Tree operating mode is set to One Spanning Tree Per VLAN (available only on AOS switch platforms).
- **Bridge ID** - The bridge identifier for this Spanning Tree instance. Consists of the bridge priority value (in hex) concatenated with the bridge MAC address.
- **Root ID** - The bridge identifier of the root of the spanning tree as determined by the Spanning Tree Algorithm and Protocol.
- **Loopback** - A loopback interface configured for testing.
- **Time Since Last Topology Change** - The amount of time, in hundredths of a second, since the last topology change was detected by this Spanning Tree instance.
- **Total Topology Change** - The number of topology changes detected by this spanning tree instance since the management entity was last reset or initialized.
- **Root Port** - The port that offers the lowest cost path from this bridge to the root bridge for this Spanning Tree instance.
- **Next Best Root Cost** - The cost of the next best root port for this Spanning Tree instance.
- **Next Best Root Port** - The port that offers the next best (second lowest) cost path to the root bridge for this Spanning Tree instance.
- **Network Maximum Age** - The Maximum Age time value for the root bridge.
- **Network Hello Time** - The Hello Time value for the root bridge.
- **Network Hold Time** - The amount of time, in hundredths of a second, in which this spanning tree instance can transmit no more than two Configuration Bridge Protocol Data Units (BPDU).
- **Network Forward Delay** - The forward delay time value for the root bridge.

Bridge View

The Bridge View List displays a Spanning Tree bridge information for devices in the selected VLAN.

- **Device Friendly Name** - The user-defined name for the device.
- **Device IP** - The IP host address that identifies the switch within the management IP network. It is not the same IP host address defined as a virtual IP router port for the selected VLAN, unless the VLAN is part of the management network. In addition, an IP host address appears in this field even if an IP virtual router port does not exist for the VLAN on that particular switch.
- **Protocol** - The VLAN spanning tree algorithm protocol. The algorithm determines the state and role of a port within the spanning tree topology:
 - **STP (802.1D)** - Standard Spanning Tree Algorithm and Protocol (Default).
 - **RSTP (802.1W)** - Rapid Spanning Tree Algorithm and Protocol. RSTP is based on the 802.1D standard algorithm and is designed to provide quick recovery in the event of a link, port or device failure.
 - **MSTP (802.1S)** - Multiple Spanning Tree Protocol. MSTP is an enhancement to 802.1Q

Common Spanning Tree Instance (CST). When the switch is running in Flat Mode, a single Spanning Tree instance is applied across all VLAN port connections. MSTP allows the configuration of Multiple Spanning Tree Instances (MSTI) 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. Note that MSTP in VLAN spanning tree view is only displayed for Instance 0 under VLAN 1. None of the other instances will be displayed.

Priority - The bridge priority value (0 - 65535) for the VLAN. The lower the number, the higher the priority value. The bridge priority value is used by the spanning tree algorithm to determine which VLAN should serve as the root of the spanning tree. (Default = 32768)

Maximum Age - The amount of time, in seconds, that spanning tree information learned from BPDUs received on VLAN ports is retained. When this information has aged beyond the maximum age value, the information is discarded. (Range = (6- 40, Default = 20)

Hello Time -The Hello Time value for the root bridge.

Forward Delay -The forward delay time value for the root bridge.

Port View

The Port View List displays a Spanning Tree port information for devices in the selected VLAN.


Device Friendly Name - The user-defined name for the device.

- **Device IP** - The IP host address that identifies the switch within the management IP network. It is not the same IP host address defined as a virtual IP router port for the selected VLAN, unless the VLAN is part of the management network. In addition, an IP host address appears in this field even if an IP virtual router port does not exist for the VLAN on that particular switch.
- **Port** - The slot/port number.
- **Priority** -The port priority value for the VLAN. The lower the number, the higher the priority value. The port priority is used to determine which port offers the best path to the root when multiple paths have the same path cost. If the priority values are the same for all ports in the path, then the port with the lowest physical switch port number is selected. (Range = (0 - 15, Default = 7)
- **Oper Status** - The operational state of the port as determined by the spanning tree algorithm:
 - **Disabled** - Physical port is down or administratively disabled.
 - **Blocking or Discarding** - Port does not transmit or receive data to prevent a network loop.
 - **Listening** - Port is preparing to transmit data.
 - **Learning** - Port is learning MAC addresses seen on the port.
 - **Forwarding** - Port is transmitting and receiving data.
- **Admin Status** - The Spanning Tree status for the port (Enabled/Disabled). If disabled, the port state is set to forwarding for the VLAN Spanning Tree instance. This status value, however, is ignored if Spanning Tree is disabled for the associated VLAN. By default, Spanning Tree is enabled on all switch ports.
- **Path Cost** - The path cost value for the port. This value specifies the contribution of a port to the path cost towards the root bridge that includes the port. If the path cost is set to 0, then a default value based on link speed is used. (Range = (0 - 65535).
- **Manual Mode** - The mode used for managing the port's state:
 - **Blocking or Forwarding (Manually Set)** - If the port state is manually set to Blocking or Forwarding, the port remains in that state until it is changed and does not participate in the spanning tree algorithm.
 - **No (Dynamic)** - Dynamic mode defers configuration of the port state to the spanning tree algorithm. By default, this parameter is set to No (Dynamic).
- **Admin Connection** - The port's administratively set connection type. This parameter is used by the 802.1w Rapid Spanning Tree Protocol (RSTP) to determine if a port is eligible for rapid transition to

the forwarding state.

- **No Point to Point** - Port connects to multiple switches.
- **Point to Point** - Port connects directly to another switch.
- **Auto Point to Point** - Connection type is automatically defined to No Point to Point or Point to Point based on the port's operational status. (Default)
- **Edge Port** - Port is at the edge of a bridged LAN, does not receive BPDU, and has only one MAC address learned. Edge ports, however, will operationally revert to a No Point to Point connection type if a BPDU is received on the port.
- **Port Role** - The role of the port for this Spanning Tree instance (e.g., Root, Designated, Alternate, Backup).
- **Oper Connection** - The operational connection type for the port:
 - **No Point to Point** - Port connects to multiple switches.
 - **Point to Point** - Port connects directly to another switch.
 - **Auto Point to Point** - Connection type is automatically defined to No Point to Point or Point to Point based on the port's operational status. (Default)
 - **Edge Port** - Port is at the edge of a bridged LAN, does not receive BPDU, and has only one MAC address learned. Edge ports, however, will operationally revert to a No Point to Point connection type if a BPDU is received on the port.

Editing an STP Configuration

You can edit Spanning Tree Bridge or Port parameters on switches in a VLAN. Select the VLAN in the VLANs Table, then select Spanning Tree from the **Actions** drop-down menu at the top of the screen to bring up the Spanning Tree Summary View Table for the VLAN. Click on the Bridge or Port link at the top of the table to bring up the Bridge or Port Table. Select a switch/port in a table then click on the Edit icon . Edit the Bridge or Port parameters as described below and click on the **Apply** button. Repeat to edit additional Bridge or Port parameters.

Note: Spanning Tree software is active on all switches by default. As a result, a loop-free network topology is automatically calculated based on default Spanning Tree switch, VLAN (bridge), and port parameter values. It is only necessary to configure Spanning Tree bridge parameters to change how the topology is calculated and maintained.

STP Bridge Parameters

- **Protocol** - The VLAN spanning tree algorithm protocol. The algorithm determines the state and role of a port within the spanning tree topology:
 - **STP (802.1D)** - Standard Spanning Tree Algorithm and Protocol (Default).
 - **RSTP (802.1W)** - Rapid Spanning Tree Algorithm and Protocol. RSTP is based on the 802.1D standard algorithm and is designed to provide quick recovery in the event of a link, port or device failure.
 - **MSTP (802.1S)** - Multiple Spanning Tree Protocol. MSTP is an enhancement to 802.1Q Common Spanning Tree Instance (CST). When the switch is running in Flat Mode, a single Spanning Tree instance is applied across all VLAN port connections. MSTP allows the configuration of Multiple Spanning Tree Instances (MSTI) 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. Note that MSTP in VLAN spanning tree view is only displayed for Instance 0 under VLAN 1. None of the other instances will be displayed.
- **Priority** - The bridge priority value (0 - 65535) for the VLAN. The lower the number, the higher the priority value. The bridge priority value is used by the spanning tree algorithm to determine which

VLAN should serve as the root of the spanning tree. (Default = 32768)

- **Maximum Age** - The amount of time, in seconds, that spanning tree information learned from BPDUs received on VLAN ports is retained. When this information has aged beyond the maximum age value, the information is discarded. (Range = (6- 40, Default = 20)
- **Hello Time** -The Hello Time value for the root bridge.
- **Forward Delay** -The forward delay time value for the root bridge.

STP Port Parameters

- **Priority** -The port priority value for the VLAN. The lower the number, the higher the priority value. The port priority is used to determine which port offers the best path to the root when multiple paths have the same path cost. If the priority values are the same for all ports in the path, then the port with the lowest physical switch port number is selected. (Range = (0 - 15, Default = 7)
- **Admin Status** - The Spanning Tree status for the port (Enabled/Disabled). If disabled, the port state is set to forwarding for the VLAN Spanning Tree instance. This status value, however, is ignored if Spanning Tree is disabled for the associated VLAN. By default, Spanning Tree is enabled on all switch ports.
- **Path Cost** - The path cost value for the port. This value specifies the contribution of a port to the path cost towards the root bridge that includes the port. If the path cost is set to 0, then a default value based on link speed is used. (Range = (0 - 65535).
- **Admin Connection** - The port's administratively set connection type. This parameter is used by the 802.1w Rapid Spanning Tree Protocol (RSTP) to determine if a port is eligible for rapid transition to the forwarding state.
 - **No Point to Point** - Port connects to multiple switches.
 - **Point to Point** - Port connects directly to another switch.
 - **Auto Point to Point** - Connection type is automatically defined to No Point to Point or Point to Point based on the port's operational status. (Default)
 - **Edge Port** - Port is at the edge of a bridged LAN, does not receive BPDU, and has only one MAC address learned. Edge ports, however, will operationally revert to a No Point to Point connection type if a BPDU is received on the port.

IP Router

The VLAN Manager IP Router Screen displays all configured IP router interfaces for the selected VLAN. It is also used to create, edit, and delete IP router interfaces on devices in the VLAN. Network device traffic is bridged (switched) at the Layer 2 level between ports that are assigned to the same VLAN. However, if a device needs to communicate with another device that belongs to a different VLAN, you must configure an IP router interface on a device in the VLAN to enable Layer 3 routing to transmit traffic between VLANs. A VLAN is available for routing when at least one IP router interface is defined for that VLAN and at least one active port is associated with the VLAN. If a VLAN does not have a router interface, the ports associated with that VLAN are in essence firewalled from other VLANs.


Creating an IP Router Interface

Click on the Create icon **+** to create an IP router interface on a device in a VLAN. Complete the fields as described below, then click on the **Create** button.

- **Device** - Select an option from the drop-down menu (Use Switch Picker/Use Topology) and click on the **Select Device** button. A list of devices that are members of the VLAN is displayed. Select the device on which you want to create an IP router interface and click **OK**.

- **Router IP Address** - The IP address of the IP router interface. Router interface IP addresses must be unique. You cannot have two router interfaces with the same IP address.
- **Router IP Mask** - The IP router subnet mask value. The default value for this field is based on which network class range the IP address falls within; Class A, B, or C. (255.0.0.0, 255.255.0.0, or 255.255.255.0).
- **IP Encapsulation** - The IP router interface frame encapsulation value (Ethernet 2 or SNAP). The frame encapsulation determines the framing type the router interface uses when generating frames that are forwarded out VLAN ports. Select an encapsulation that matches the encapsulation of the majority of IP VLAN traffic. (Default = Ethernet 2)
- **IP Forwarding** - The router interface forwarding status (Enabled/Disabled). A forwarding router interface sends IP frames to other subnets. A "no forwarding" router interface acts as a host only. It receives IP frames from other router interfaces. (Default = Enabled).
- **Interface Name** - The user-defined interface name (up to 20 characters).
- **VRF ID** - The VRF ID. If configured, select a VRF from the drop-down menu to assign the interface to a configured VRF instance (by default all interfaces are assigned to the Default VRF). You can assign a new interface to a VRF; however, you cannot edit the VRF ID of an existing interface. If the feature is not available on the device, the column will display "Default", indicating that the switch is operating as a single routing instance. VRF instances are created on the switch through the CLI or WebView application. Click here for more information on the Multiple VRF Feature. See the "Configuring Multiple VRF" Chapter in the applicable *OmniSwitch AOS Network Configuration Guide* for detailed instructions on configuring VRF instances.

Editing an IP Router Interface

To edit an IP router interface on a device, select the interface in the IP Router Table and click on the Edit icon . Edit any fields as described above and click on the **Apply** button. Note that you cannot edit the Interface Name or VRF fields.

Deleting an IP Router Interface

To delete an IP router interface, select the interface in the IP Router Table, click on the Delete icon, then click **OK** at the Confirmation Prompt.

Viewing IP Router Interfaces

The Router Screen displays all configured IP router interfaces for the selected VLAN. The fields are defined below.

- **Device Name** - The user-configured name for the device on which the interface is configured.
- **Device Friendly Name** - The IP address of the device.
- **Device IP Address** - The IP address of the device.
- **Router IP Address** - The IP address of the IP router interface.
- **Router IP Mask** - The IP router subnet mask value. The default value for this field is based on which network class range the IP address falls within; Class A, B, or C. (255.0.0.0, 255.255.0.0, or 255.255.255.0).
- **IP Encapsulation** - The IP router interface frame encapsulation value (Ethernet 2 or SNAP). The frame encapsulation determines the framing type the router interface uses when generating frames that are forwarded out VLAN ports. Select an encapsulation that matches the encapsulation of the majority of IP VLAN traffic. (Default = Ethernet 2)
- **IP Forwarding** - The router interface forwarding status (Enabled/Disabled). A forwarding router interface sends IP frames to other subnets. A "no forwarding" router interface acts as a host only. It receives IP frames from other router interfaces. (Default = Enabled).

- **Interface Name** - The user-defined interface name (up to 20 characters).
- **VRF ID** - The VRF ID. If multiple VRFs are configured on the device, the VRF ID is displayed. If none are configured, or if the feature is not available on the device, the column will display "Default", indicating that the switch is operating as a single routing instance.

VLAN Details

The VLAN Details Screens provide detailed information about devices in the selected VLAN. By default, the Device View is displayed. Click on a link at the top of the screen to display Device, Port, or Link Aggregate information.

Device View

- **Friendly Name** - The IP address of the device.
- **Device Name** - The user-configured name for the device.
- **Device IP Address** - The IP address of the device.
- **Device MAC Address** - The MAC address of the device.
- **Device Version** - The version number of the device software. OmniVista may not be able to determine the software version on some third-party devices. In these cases, the field will be blank.
- **Device Location** - The physical location of the device (user-defined, up to 255 characters). If the user did not specify a location the field displays "Unknown".
- **Device Status** - This field displays the operational status of the
 - device: Up - Device is up and responding to polls.
 - Down - Device is down and not responding to polls.
 - Warning - Device has sent at least one warning or critical trap
- **Device Type** - The device model (e.g., OS6900-X72).
- **VLAN ID** - The VLAN ID number (e.g., VLAN 10).
- **VLAN Description** - A text string up to 32 characters. This parameter defaults to the VLAN ID number (e.g., VLAN 10) if a description is not specified at the time the VLAN is created.
- **VLAN Admin Status** - The administrative status of the VLAN (Enabled/Disabled). By default, the administrative status is enabled when a VLAN is created. When a VLAN is administratively disabled, static port and dynamic mobile port assignments are retained but traffic on these ports is not forwarded. However, VLAN rules remain active and continue to classify mobile port traffic for VLAN membership.
- **VLAN Type** - The type of VLAN is determined at the time the VLAN is created (e.g., Standard, BVLAN, Control BVLAN).
- **Spanning Tree Status** - The Spanning Tree Status for the VLAN (Enabled/Disabled). When a VLAN is created, an 802.1D standard Spanning Tree Algorithm and Protocol (STP) instance is enabled for the VLAN by default. STP evaluates VLAN port connections to determine if there are redundant data paths between the same VLAN on other switches. If a redundant path does exist, STP determines which path to block in order to provide a loop-free network topology. In this manner, STP ensures that there is always only one active data path between any two switches (VLANs). When a change occurs, such as a path is disconnected or a path cost change, the Spanning Tree Algorithm activates the blocked path to restore the network connection.
- **Mobility** - The mobile status for the VLAN (Enabled/Disabled). On AOS switches, mobility is not enabled or disabled at the VLAN level. Instead, switch ports are designated as mobile or non-mobile.
- **Oper Status** - The VLAN operational status (Active/Inactive). This parameter is not modifiable; switch software determines if the VLAN is operationally active or inactive and sets the appropriate field value. A VLAN's operational status remains inactive until at least one active switch port is assigned to the VLAN and the VLAN's administrative status is enabled. This means that VLAN properties, such as

Spanning Tree or router ports, also remain inactive. Ports are considered active if they are connected to an active network device. Non-active port assignments are allowed, but do not change the VLAN's operational state.

- **Authentication** - The authentication status for the VLAN (Enabled/Disabled). By default, authentication is disabled when a VLAN is created. Once authentication is enabled on a VLAN, however, then only authenticated mobile port devices can join the VLAN after completing the appropriate log-in process. Layer 2 authentication uses VLAN membership to grant access to network resources. Authenticated VLANs control membership through a log-in process; this is sometimes called user authentication. A VLAN must have authentication enabled before it can participate in the Layer 2 authentication process.
- **Voice Status** - Administrative status (Enabled/Disabled) of voice usage for the current VLAN (supported on 6.x devices only).
- **Router Protocol** - The protocol for the VLAN virtual router port. If no router port is configured for the VLAN, then "N/A" appears in this field. A VLAN is available for routing when a virtual router port is defined for that VLAN and at least one active port has joined the VLAN. If a VLAN does not have a router port, its ports are in essence firewalled from other VLANs.

Port View

- **Device Friendly Name** - The IP address of the device.
- **Device IP Address** - The IP address of the device.
- **Port** - The VLAN slot/port number.
- **Port Type** - The Port Type parameter indicates how the port assignment to the current VLAN was made.
 - **Default** - The port is a fixed port that was statically assigned to the VLAN, which is now the configured default VLAN for the port.
 - **Q Tagged** - The port is a fixed port that was statically assigned to the VLAN using the 802.1Q tagging feature. The VLAN is a static secondary VLAN assignment for the 802.1Q tagged port.
 - **Mobile** - The port is a mobile port that was dynamically assigned to the VLAN when traffic received on the port match traffic rules defined for the VLAN. The VLAN is a dynamic secondary VLAN assignment for the mobile port.
- **Port State** - The status of the VLAN port assignment.
 - **Forwarding** - Port is active and forwarding traffic on this VLAN.
 - **Inactive** - Port is not active (administratively disabled, down, or nothing is connected to the port).
 - **Blocking** - Port is active, but not forwarding any traffic on this VLAN.
 - **Filtering** - Mobile port traffic is filtered for the VLAN; only traffic received on the port that matches VLAN rules is forwarded. Occurs when a mobile port's VLAN is administratively disabled or the port's default VLAN status is disabled. Does not apply to fixed ports.

Link Aggregate View



- **Friendly Name** - The IP address of the device.
- **Device IP Address** - The IP address of the device.
- **Link Aggregate ID** - The ID of the link aggregate group of ports. This number was assigned when the aggregate was created.
- **Link Aggregate** - An optional textual description (up to 32 characters) for the link aggregate.
- **Port Type** - The Port Type parameter indicates how the port assignment to the current VLAN was made.
 - **Default** - The port is a fixed port that was statically assigned to the VLAN, which is now

the configured default VLAN for the port.

- **Q Tagged** - The port is a fixed port that was statically assigned to the VLAN using the 802.1Q tagging feature. The VLAN is a static secondary VLAN assignment for the 802.1Q tagged
- port. **Mobile** - The port is a mobile port that was dynamically assigned to the VLAN when traffic received on the port match traffic rules defined for the VLAN. The VLAN is a dynamic secondary VLAN assignment for the mobile port.
- **Description** - The standard MIB name for this dynamic aggregate group.

Device Selection

The Device Selection Screen is used to configure basic VLAN configuration parameters and select devices to be included in the VLAN. You can create a single VLAN or create multiple VLANs. When you have completed all of the parameters, click the **Next** button at the bottom of the screen or click on VLAN Configuration on the left side of the screen to move to the next step.

- **VLAN ID** - The VLAN ID number (Range = 2 - 4094). To create a single VLAN, enter a single VLAN ID and click on the Add icon . Repeat to create multiple VLANs, or enter a range of VLANs (e.g., 500 - 505) and click on the Add icon . Note that you cannot enter a VLAN ID for a VLAN that has already been created. If you want to add additional devices to an existing VLAN, return to the VLANs Table and edit the VLAN to add the additional devices.
- **Description** - Enter an optional description for the VLAN. If you do not enter a description, the VLAN ID is used.
- **VLAN Configuration Template** - The VLAN Configuration Template contents are applied to all selected devices, subject to availability of ports on devices. Users can override these settings.
 - **Admin Status** - The administrative status of the VLAN (Enabled/Disabled).
 - **Default VLAN ID** (Multiple VLAN Configuration Only) - If you are configuring multiple VLANs, enter one of the VLAN ID of one of the VLANs. This will be the default VLAN for the configured ports.
 - **Default Ports Template** - Used to configure default (untagged) ports for all selected devices. The configuration will be added to selected devices if the port is available on device. For example , if you enter port 1/5 as a default port when creating VLAN 100, port 1/5 will be added as a default port for all devices having port 1/5. If a selected device does not have port 1/5, the default port will not be created on the device. Ports or Link Aggregates must be entered in the format shown (e.g., LAG-1, lag-1, 1/1, 1/2a, 1/1-1/7, 1/1/1, 1/1/1a, 1/1/1-1/1/7). Note that you will be able to add additional ports (or remove ports) on the Default Ports Assignment Screen later in the wizard.
 - **Q Tagged Ports Template** - Used to configure Q-Tagged ports for all selected devices. The configuration will be added to selected devices if the port is available on device. For example , if you enter port 1/3 as a Q-Tagged Port when creating VLAN 100, port 1/3 will be added as a Q- Tagged port for all devices having port 1/3. If a selected device does not have port 1/3, the Q- Tagged port will not be created on the device. Ports or Link Aggregates must be entered in the format shown (e.g., LAG-1, lag-1, 1/1, 1/2a, 1/1-1/7, 1/1/1, 1/1/1a, 1/1/1-1/1/7). Note that you will be able to add additional ports (or remove ports) on the Q-Tagged Ports Assignment Screen later in the wizard. Also note that for wireless devices, Q-Tagged Port configuration is supported on trunk ports.
 - **Device Selection** - Select an option from the drop-down menu (Use Switch Picker/ Use Topology) and click **Add Remove Device** button to select devices for the VLAN.

Note: When creating a VLAN, you can select up to 200 devices. If necessary, once the VLAN is created you can edit the VLAN to add additional devices. Again, for each edit, you can add up to 200 devices. Repeat to add additional devices.

VLAN Configuration

The VLAN Configuration Screen is used to review VLAN device selection and review/modify VLAN administrative status. If necessary, click the **Back** button to modify the device selection, or click on the **Admin state configuration** button to change the VLAN administrative state. When you are finished, click the **Next** button at the bottom of the screen or click on Default Ports Assignment on the left side of the screen to move to the next step.

Default Ports Assignment

The Default Ports Assignment Screen is used to configure ports on the selected device(s) to be included in the VLAN. Click on a device in the list and click on the Add Ports for Device ... button (or just click on the "Add Port" link under a device) to bring up the Port Selection Window. Select the device ports to be included in the VLAN and click OK. Repeat to add ports for additional devices.

Note that ports added on the Device Selection Screen of the wizard in the VLAN Configuration Template will be "pre-selected" on the Default Ports Selection window. You can add additional ports or remove ports. Also note that Q-Tagged Ports will not be available for selection.

After selecting ports for each device, click the **Next** button at the bottom of the screen or click on Q Tagged Ports Assignment on the left side of the screen to move to the next step.

Q Tagged Ports Assignment

The Q Tagged Ports Assignment Screen is used to configure Q-Tagged Ports on the selected device(s) to be included in the VLAN. Click on a device in the list and click on the Add Q Tagged Ports for Device ... button (or just click on the "Add Port" link under a device) to bring up the Port Selection Window. Select the device ports to be included in the VLAN and click OK. Repeat to add ports for additional devices.

Note that ports added on the Device Selection Screen of the wizard in the VLAN Configuration Template will be "pre-selected" on the Tagged Ports Selection window. You can add additional ports or remove ports. Also note that ports selected on Default Port Assignment Screen will not be available for selection.

After selecting ports for each device, click the Next button at the bottom of the screen or click on Review on the left side of the screen to move to the next step.

Review

The Review Screen is used to review the VLAN configuration. VLAN Configuration Template information is displayed at the top of the window; and a list of switches contained in the configured VLAN is displayed. By default the List of Switches displays device configuration information. Click on a link at the top of the table to display Port, or Link Aggregate information.

After reviewing the configuration, click the Create button to create the VLAN. You can also click the Back button to return to a previous screen and modify the configuration before returning to this screen to create it.

MVRP Overview

Multiple VLAN Registration Protocol (MVRP) provides a mechanism for dynamic maintenance of the contents of dynamic VLAN registration entries for each VLAN, and for propagating the information they contain to other bridges. This information allows MVRP-aware devices to dynamically establish and update their knowledge of the set of VLANs that currently have active members, and through which ports those members can be reached. The main purpose of MVRP is to allow switches to automatically discover some of the VLAN information that would otherwise have to be manually configured.

MVRP acts as a Multiple Registration Protocol (MRP) application, sending and receiving MVRP information encapsulated in an ethernet frame on a specific MAC address. MVRP allows both end stations and bridges in a bridged local area network to issue and revoke declarations relating to membership of VLANs. Each MVRP device that receives the declaration in the network creates or updates a dynamic VLAN registration entry in the filtering database to indicate that the VLAN is registered on the reception port. For more information, see the "Configuring MVRP" Chapter in the applicable *OmniSwitch Network Configuration Guide*.

The MVRP Configuration Wizard guides you through the steps to configure MVRP on network switches/ports. Note that the MVRP feature is supported only in STP flat mode. If MVRP is configured in the system with STP flat mode, STP mode cannot be changed to per-VLAN mode. When a topology change is detected by STP, MAC addresses for the dynamic VPAs learned by MVRP is also deleted.

Summary View

The Multiple VLAN Registration Protocol (MVRP) Summary View Screen displays an overview of MVRP switch and port configuration on the network. MVRP provides a mechanism for dynamic maintenance of the contents of dynamic VLAN registration entries for each VLAN, and for propagating the information they contain to other bridges. This information allows MVRP-aware devices to dynamically establish and update their knowledge of the set of VLANs that currently have active members, and through which ports those members can be reached. The main purpose of MVRP is to allow switches to automatically discover some of the VLAN information that would otherwise have to be manually configured.

MVRP acts as an MRP application, sending and receiving MVRP information encapsulated in an ethernet frame on a specific MAC address. MVRP allows both end stations and bridges in a bridged local area network to issue and revoke declarations relating to membership of VLANs. Each MVRP device that receives the declaration in the network creates or updates a dynamic VLAN registration entry in the filtering database to indicate that the VLAN is registered on the reception port.

Switches List

The Switches List provides an overview of MVRP configuration on network switches.

- **Device Friendly Name** - The IP address of the device.
- **Device IP** - The IP address of the device.
- **MVRP Status** - The MVRP administrative status on the switch (Enabled/Disabled)
- **Transparent Switch** - The administrative status of MVRP transparent switching on the switch. If enabled, when MVRP is globally disabled on the device, MVRP frames are flooded transparently. If disabled, the device will discard received MVRP frames. (Default = Disabled)
- **Max VLAN Limit** - The maximum number of dynamic VLANs that can be created by MVRP on the switch. (Range = 32 - 4094, Default = 256)
- **Registration Protocol** - The registration protocol running on the switch (MVRP/GVRP). (Default = MVRP)

Ports Configuration

Click on a switch in the Switches List to view MVRP port information for the switch.

- **Port** - The slot/port number of the MVRP interface. MVRP can be enabled on switch ports regardless of whether it is globally enabled on the switch. However, for the port to become an active participant in the MVRP operation, MVRP must be enabled globally on the switch. When MVRP is globally enabled on the switch and is not enabled on the port, that port is excluded from the MVRP protocol operation.
- **MVRP Status** - MVRP status for the port (Enabled/Disabled). MVRP can be enabled on switch ports regardless of whether it is globally enabled on the switch. However, for the port to become an active participant in the MVRP operation, MVRP must be enabled globally on the switch. When MVRP is globally enabled on the switch and is not enabled on the port, that port is excluded from the MVRP protocol operation. Note: MVRP can be enabled only on fixed ports, 802.1 Q ports, aggregate ports, and VLAN Stacking Network ports. Other ports (e.g., mirroring ports, VLAN Stacking User ports) do not support MVRP.
- **Registrar Mode** - The MVRP Registration Mode of the port:
 - **Normal** - Specifies that both registration and de-registration of VLANs is allowed. VLANs can be mapped either dynamically (through MVRP) or statically (through management application). (Default)
 - **Fixed** - Specifies that only static mapping of VLANs is allowed on the port, but de-registration of previously created dynamic or static VLANs is not allowed.
 - **Forbidden** - Specifies that dynamic VLAN registration or de-registration is not allowed on the port. Any dynamic VLANs created earlier are de-registered.
- **Applicant Mode** - The Applicant Mode of the port. This configures whether MVRP PDU exchanges are allowed on the port, depending on the port's Spanning Tree state:
 - **Participant** - MVRP PDU exchanges are only allowed when the port is in the STP forwarding state.
 - **Non-Participant** - MVRP PDUs are not sent in this mode, and PDUs received are processed as expected.
 - **Active** - MVRP PDU exchanges are allowed when the port is in the STP forwarding state or STP blocking state.
- **Periodic Timer** - The MVRP periodic-timer time interval, in seconds, for dynamically registering VLANs on the switch. The default timer setting is used unless there is a compelling reason to change the settings. Modifying timers to inappropriate values can cause an imbalance in the operation of MVRP. (Default = 1).
- **Periodic Transmission** - The periodic transmission status on the port (Enabled/Disabled). (Default = Disabled)

Global Parameters

The MVRP Configuration Wizard Global Parameters Screen is used to configure global MVRP parameters on a switch. You can apply the existing global configuration to switches by selecting the **Keep Existing Global Configuration** radio button or you can select the **Apply New Configuration** radio button and change the global MVRP parameters as described below. When you are finished, click on the **Next** button or select **Port Parameters** on the left side of the screen to move to the next screen.

Note: To view the current global configuration, click on the **Apply New Configuration** radio button.

Global Parameter Fields

- **MVRP Status** - Enables/Disables MVRP globally on the switch. To enable MVRP on a port, MVRP must be enabled on the switch. Disabling MVRP globally deletes all VLANs learned through MVRP. MVRP is supported only when the switch is operating in the flat Spanning Tree mode.
- **Transparent Switching** - Enables/Disables transparent switching for MVRP. If enabled, when MVRP is globally disabled on the device, MVRP frames are flooded transparently. If disabled, the device will discard received MVRP frames. (Default = Disabled)
- **Max VLAN Limit** - The maximum number of dynamic VLANs that can be created by MVRP. The Max VLAN Limit can be configured even if MVRP is not enabled on the switch. However, MVRP must be enabled on the switch for creating dynamic VLANs. If you set the VLAN limit to less than the current number of dynamically learned VLANs, the new configuration takes effect only after MVRP is disabled and re-enabled on the switch. The VLANs learned earlier are retained if this operation is not performed. (Range = 32 - 4094, Default = 256)
- **Registration Protocol** - The registration protocol running on the switch (MVRP/GVRP). (Default = MVRP)

Port Parameters

The MVRP Configuration Wizard Port Parameters Screen is used to configure global MVRP port parameters. You can apply the existing MVRP port parameters to switches by selecting the **Keep Existing Port Configuration** radio button or you can select the **Apply New Configuration** radio button and change the MVRP port parameters as described below. When you are finished, click on the **Next** button or select **Devices/Ports** on the left side of the screen to move to the next screen.

Note: To view the current MVRP port parameters, click on the **Apply New Configuration** radio button.

Port Parameter Fields

- **MVRP Status** - The MVRP port status (Enabled/Disabled). MVRP can be enabled on switch ports regardless of whether it is globally enabled on the switch. However, for the port to become an active participant in the MVRP operation, MVRP must be enabled globally on the switch. When MVRP is globally enabled on the switch and is not enabled on the port, that port is excluded from the MVRP protocol operation. Note that MVRP can only be enabled on fixed ports, 802.1 Q ports, aggregate ports, and VLAN Stacking Network ports. Other ports (e.g., mirroring ports, VLAN Stacking User ports) do not support MVRP.
- **Registrar Mode** - The MVRP Registration Mode of the port:
 - **Normal** - Both registration and de-registration of VLANs is allowed. VLANs can be mapped either dynamically (through MVRP) or statically (through management application). (Default)
 - **Fixed** - Only static mapping of VLANs is allowed on the port, but de-registration of previously created dynamic or static VLANs is not allowed.
 - **Forbidden** - Dynamic VLAN registration or de-registration is not allowed on the port. Any dynamic VLANs created earlier are de-registered.
- **Applicant Mode** - The Applicant Mode of the port. This configures whether MVRP PDU exchanges are allowed on the port, depending on the port's Spanning Tree state:
 - **Participant** - MVRP PDU exchanges are only allowed when the port is in the STP forwarding state.
 - **Non-Participant** - MVRP PDUs are not sent in this mode, and PDUs received are processed as expected.

- **Active** - MVRP PDU exchanges are allowed when the port is in the STP forwarding state or STP blocking state. (Default)
- **Periodic Timer** - The MVRP periodic-timer time interval, in seconds, for dynamically registering VLANs on the switch. Use default timer setting unless there is a compelling reason to change the settings.
Modifying timers to inappropriate values can cause an imbalance in the operation of MVRP. (Default = 1)
- **Periodic Transmission** - Enables/Disables the periodic transmission status on the port. (Default = Disabled)

Devices/Ports

The MVRP Configuration Wizard Devices/Ports Screen is used to apply the MVRP global/port parameters to specific switches/ports). In the Devices area, select an option from the drop-down menu to select switches, then click on the Add/Remove Devices button. The selected switches will be listed. Select a switch in the list and click on the Add/Remove Ports button to select specific ports on the switch. Repeat the process for additional switches.

When you are finished, click on the **Next** button or select **Review** on the left side of the screen to move to the next screen.

Review

The MVRP Configuration Wizard Review Screen is used to review your MVRP configuration before applying it to the selected switches/ports. If necessary, click on the **Back** button to return to a screen to make any changes to the configuration. When you are finished, click on the **Apply** button to apply the configuration.

The operation will be displayed on the Results Screen. Click **OK** to return to the top of the MVRP Configuration Wizard.

IP Interface

The VLAN Manager IP Interface Screen is used to view configured IP interfaces on a device; and to create, edit, and delete IP interfaces on devices. You can configure up to eight (8) IP interfaces per VLAN on OS6800/6850/9000/9000E Switches (Release 6.1.1 and later); and up to sixteen (16) IP interfaces per VLAN on OS6860, OS6900 (Release 7.2.1 and later), OS9900, and OSOS10K Switches (Release 7.1.1 and later).


Creating an IP Interface

To create an IP interface on a device, click on the Create icon **+**, complete the fields as described below, then click on the **Create** button.


- **Name** - The user-configured interface name.
- **IP Address** - The IP address of the IP interface.
- **Subnet Mask** - The IP interface subnet mask value. The default value for this field is based on which network class range the IP address falls within; Class A, B, or C. (255.0.0.0, 255.255.0.0, or 255.255.255.0).
- **MTU** - The Maximum Transmission Unit size set for the interface.
- **Device Type** - The type of device bound to the interface:
 - **Unbound** - No device is bound to the interface.
 - **VLAN** - Associates a VLAN with the interface. You must enter the VLAN ID in the VLAN ID field.

- **EMP** - The Ethernet Management Port is bound to the interface.
- **Loopback** - A loopback interface configured for testing.
- **GRE Tunnel** - A GRE Tunnel is configured for the interface. You must enter the Tunnel Source and Destination.
- **IPIP Tunnel** - An IPIP Tunnel is configured for the interface. You must enter the Tunnel Source and Destination.
- **VRF** - The VRF ID. If the device supports the Multiple VRF feature, select a VRF from the drop-down menu to assign the interface to a configured VRF instance (by default all interfaces are assigned to the Default VRF). You can assign a new interface to a VRF; however, you cannot edit the VRF ID of an existing interface. If the feature is not available on the device, the column will display "Default", indicating that the switch is operating as a single routing instance. VRF instances are created on the switch through the CLI or WebView application. Click here for more information on the Multiple VRF Feature. See the "Configuring Multiple VRF" Chapter in the applicable *OmniSwitch AOS Network Configuration Guide* for detailed instructions on configuring VRF instances.
- **IP Encapsulation** - The IP interface frame encapsulation value (Ethernet 2 or SNAP). The frame encapsulation determines the framing type the interface uses when generating frames that are forwarded out VLAN ports. Select an encapsulation that matches the encapsulation of the majority of IP VLAN traffic. (Default = Ethernet 2)
- **Admin State** - The administrative state of the interface (Enable/Disabled). (Default = Enabled)
- **IP Forward** - The interface forwarding status (Enabled/Disabled). A forwarding interface sends IP frames to other subnets. A "no forwarding" interface acts as a host only. It receives IP frames from other interfaces. (Default = Enabled).
- **Local Proxy ARP** - Local Proxy ARP status (Enabled/Disabled). The Local Proxy ARP feature is an extension of the Proxy ARP feature, but is enabled on an IP interface and applies to the VLAN bound to that interface. When Local Proxy ARP is enabled, all ARP requests received on VLAN member ports are answered with the MAC address of the IP interface that has Local Proxy ARP enabled. In essence, all VLAN traffic is now routed within the VLAN instead of bridged. (Default = Disabled)
- **Primary Interface** - If set to "True", designates the IP interface as the primary interface for the VLAN. If set to "False", the first interface bound to the VLAN becomes the primary by default. (Default = False)
- **Devices** - Use the Switch Picker or Topology application to select the device on which you want to configure the interface.

Editing an IP Interface

To edit an IP interface, select the interface in the IP Interface List and click on the Edit icon . Edit any fields as described above and click on the **Apply** button. Note that you cannot edit the Interface Name, Device Type, or VRF fields.

Deleting an IP Interface

To delete an IP interface, select the interface in the IP Interface List, click on the Delete icon , then click **OK** at the Confirmation Prompt.

Viewing IP Interfaces

The Interface List displays basic information for all configured IP interfaces for a selected device. To view IP interfaces on a device, select an option from the drop-down menu (Use Switch Picker/Use Topology) and click on the **Select Device** button. Select a device and click **OK**. Information for all IP interfaces configured on the device is displayed in the IP Interface List as described below. Click on an entry in the list to view detailed information about an interface.

Basic Information

- **Name** - The user-configured interface name.
- **IP Address** - The IP address of the IP interface.
- **Subnet Mask** - The IP interface subnet mask value.
- **Admin State** - The administrative state of the interface (Enable/Disabled).
- **Oper State** - The operational state of the interface.
- **Device Type** - The type of device bound to the interface:
 - **Unbound** - No device is bound to the interface.
 - **VLAN** - Associates a VLAN with the interface. You must enter the VLAN ID in the VLAN ID field.
 - **EMP** - The Ethernet Management Port is bound to the interface.
 - **Loopback** - A loopback interface configured for testing.
 - **GRE Tunnel** - A GRE Tunnel is configured for the interface. You must enter the Tunnel Source and Destination. The GRE Tunnel devices are supported only on OS10K Switches.
- **IP Forward** - The interface forwarding status (Enabled/Disabled). A forwarding interface sends IP frames to other subnets. A "no forwarding" interface acts as a host only. It receives IP frames from other interfaces. (Default = Enabled).
- **MTU** - The Maximum Transmission Unit size set for the interface.

Detailed Information

- **Name** - The user-configured interface name.
- **IP Address** - The IP address of the IP interface.
- **Subnet Mask** - The IP interface subnet mask value. The default value for this field is based on which network class range the IP address falls within; Class A, B, or C. (255.0.0.0, 255.255.0.0, or 255.255.255.0).
- **Admin State** - The administrative state of the interface (Enable/Disabled). (Default = Enabled)
- **Device Type** - The type of device bound to the interface:
 - **Unbound** - No device is bound to the interface.
 - **VLAN** - Interface is associated with a VLAN. (The VLAN ID will be displayed in the next field.)
 - **EMP** - The Ethernet Management Port is bound to the interface.
 - **Loopback** - A loopback interface configured for testing.
 - **GRE Tunnel** - A GRE Tunnel is configured for the interface. (The GRE Tunnel Source and Destination will be displayed in the next field.)
- **IP Forward** - The interface forwarding status (Enabled/Disabled). A forwarding interface sends IP frames to other subnets. A "no forwarding" interface acts as a host only. It receives IP frames from other interfaces. (Default = Enabled).
- **IP Encapsulation** - The IP interface frame encapsulation value (Ethernet 2 or SNAP). The frame encapsulation determines the framing type the interface uses when generating frames that are forwarded out VLAN ports. Select an encapsulation that matches the encapsulation of the majority of IP VLAN traffic. (Default = Ethernet 2)
- **VRF - VRF ID** - The VRF ID. If multiple VRFs are configured on the device, the VRF ID is displayed. If none are configured, or if the feature is not available on the device, the column will display "Default", indicating that the switch is operating as a single routing instance.
- **MTU** - The Maximum Transmission Unit size set for the interface.
- **Local Proxy ARP** - Local Proxy ARP status (Enabled/Disabled). (Default = Disabled)

- **Primary Interface** - If set to "True", designates the IP interface as the primary interface for the VLAN. If set to "False", the first interface bound to the VLAN becomes the primary by default.
- (Default = False)
- **Oper State** - The operational status of the router interface; Active or Inactive. An IP router interface is not operationally active until at least one active switch port is assigned to the VLAN. This is not a configurable parameter; switch software automatically determines the operational status of the VLAN and router interface.
- **Oper Reason** - An explanation of the operational state. If the interface is up the field will indicate "Interface Up". If the interface is down, an explanation is displayed:
 - **Unbound** - No device is bound to the interface.
 - **Device Down** - Device bound to the interface is down. **Admin Down** - The admin state of the interface is down. **No Such Device** - Device does not exist.
 - **No Router MAC** - No MAC address available for the interface.
 - **Tunnel Source Invalid** - The source IP address of the tunnel is invalid.
 - **Tunnel Destination Unreachable** - The destination IP address of the tunnel is not reachable.
- **Router MAC** - The switch MAC address assigned to the interface. Each interface assigned to the same VLAN shares the same switch MAC address.
- **Broadcast Address** - The default broadcast address value. The default value for this field is based on the default network class range of the IP address assigned to the router interface. For example, a class A IP address, such as 10.0.0.2, has a default broadcast address of 10.255.255.255. A class C address, such as 198.181.10.2, has a default broadcast address of 198.181.10.255.
- **Actual Primary** - Indicates if the interface is the configured and/or actual primary interface for the device (True/False).

Poll

Network devices are polled every 30 minutes for VLAN information update. The VLAN Manager Poll Screen can be used at any time to manually poll devices VLAN updates. Select an option from the drop-down menu (Use Switch Picker/User Topology), click on the **Select Devices** button, then click on the **Start Polling** button. A Result Screen will appear to indicate if the poll was successful. Click **OK** to return to the Poll Screen. VLAN information displayed in the VLAN Manager application will be updated for the selected devices.

Note: You can poll a maximum of 200 devices at a time.

Multiple Virtual Routing and Forwarding

The Multiple Virtual Routing and Forwarding (VRF) feature allows the user 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. Each VRF instance is in essence a virtual LAN for Layer 3 traffic. Some of the benefits of using the Multiple VRF feature include:

- 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 to those interfaces 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 may operate within one or more VRF instances.
- 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.

This implementation of VRF functionality does not require a BGP/MPLS configuration in the provider network. Instead, VRF instances can route and forward IP traffic between customer sites using point-to-point Layer 3 protocols, such as IP-IP or GRE tunneling.

Note: SNMPv3 is required to discover and manage VRF instances; SNMPv1 and v2 are not supported. If OmniVista finds Multiple VRFs configured on a device using SNMPv2 during Discovery polling, OmniVista warns the user about presence of Multiple VRFs and logs the message.

Configuring the Multiple VRF Feature

VRF instances are created using the Command Line Interface (CLI) or WebView application. Configuring the Multiple VRF feature consists of creating a VRF instance, assigning one or more IP interfaces to the instance, and configuring routing protocols to operate within a specific instance.

The initial configuration of an Alcatel-Lucent Enterprise switch consists of a default VRF instance, which is always active when the switch starts up and is not removable from the switch configuration. Any subsequent configuration of switch applications applies only to the default instance. To provide multiple, independent IP routing domains on the same switch, configuring additional VRF instances is required.

Creating a VRF Instance Using the CLI

Use the CLI **vrf** command to create a VRF instance. A VRF instance is identified by a name, which is specified at the time the instance is configured. For example, the following command creates the IpOne instance:

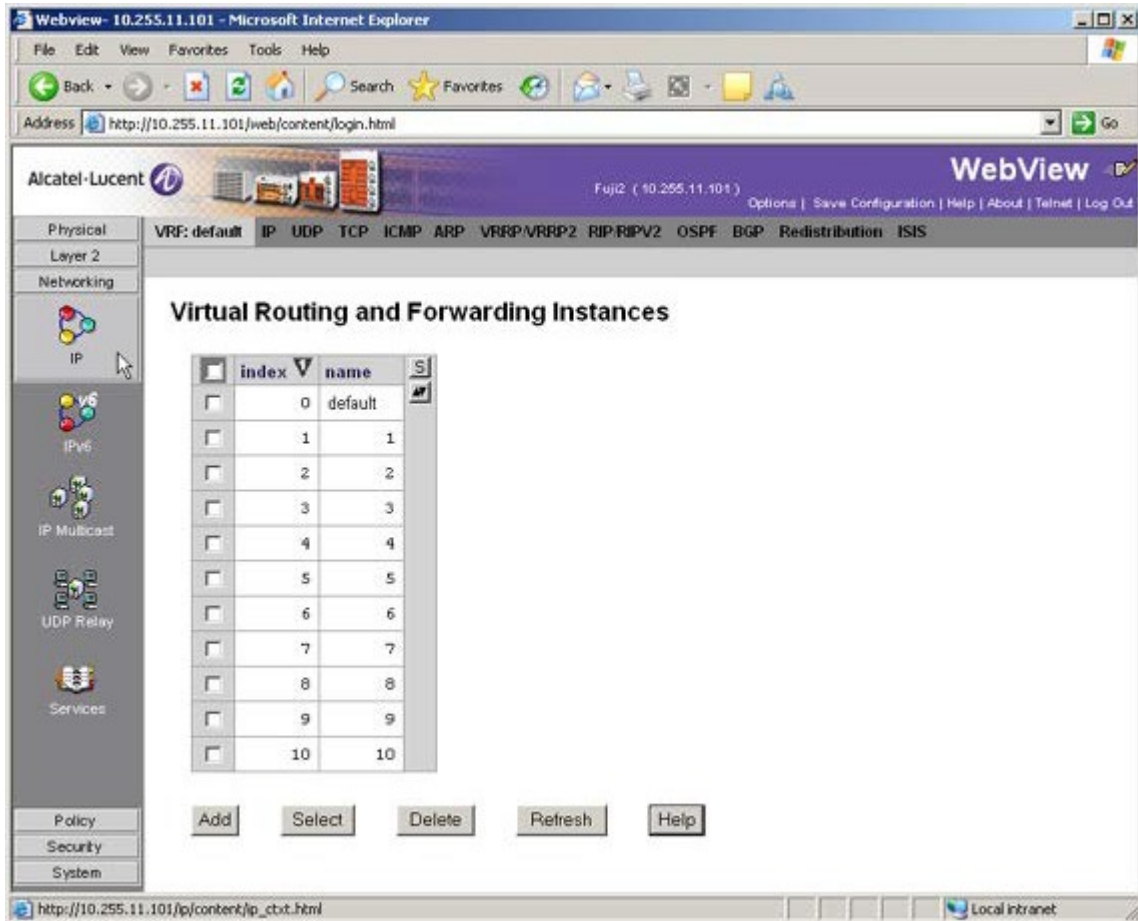
```
-> vrf IpOne
```

Use the **vrf** command to configure additional instances on the switch. Once you configure VRF instances, they will appear in the VRF ID drop-down field on the **View/New IP Router** Panels.

Note: See the "Configuring Multiple VRF" Chapter in the applicable *OmniSwitch AOS Network Configuration Guide* for detailed instructions on configuring VRF instances.

Creating a VRF Instance Using the WebView Application

Go to **Networking - IP - VRF** to bring up the Virtual Routing and Forwarding Instances Table. Use this page to create a new VRF instance or select an existing instance for configuration.



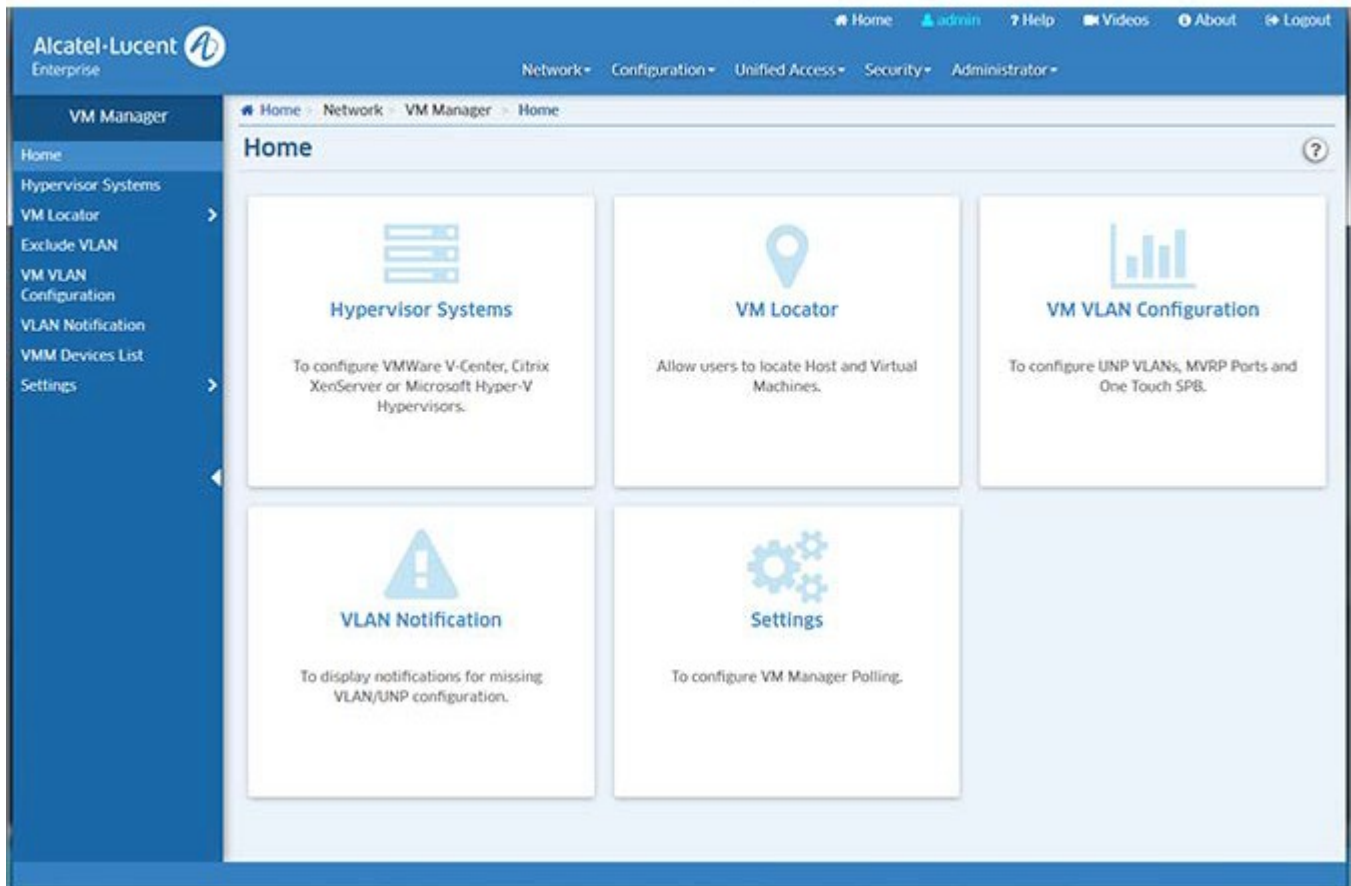
Note: See WebView Help for detailed instructions on configuring VRF instances using the WebView application.

27.0 VM Manager

Virtualization allows multiple Virtual Machines to run in isolation, side-by-side on the same physical machine (Host Server). Each virtual machine can interact independently with other devices, applications, data and users as though it were a separate physical resource. This enables much more efficient and reliable use of server resources because different Virtual Machines can run different operating systems and multiple applications while sharing the resources of a single physical machine. And because each Virtual Machine is isolated from other Virtual Machines, if one crashes, it does not affect the others. Moreover, Virtual Machines can dynamically migrate between Hosts to better utilize server resources.

Virtual Machines are configured using third-party software (VMware's vSphere, Citrix XenServer, or Microsoft Hyper-V). The OmniVista Virtual Machine (VM) Manager application interfaces with vSphere, XenServer, or Hyper-V to provide a single GUI interface to easily monitor Virtual Machines, including tracking Virtual Machines and their network associations if the machines move to a different Host on the network. Moreover, VM Manager interfaces with the Universal Network Profile (UNP) feature within OmniVista's Access Guardian application to shape Virtual Machine traffic based on user-configured UNP rules (e.g., VLAN Tag Rules, IP Rules, MAC Range Rules). Click here for an overview of Virtualization and VM Manger's role in managing Virtual Machines.

VM Manager



Note: VM Manager supports a mixture of vCenters, XenServers and Hyper-V Servers in the same configuration. You can manage up to a total of 5,000 Virtual Machines (i.e., 5,000 VMs total on all Hypervisors).

Note: The OmniVista Server can run on a Virtual Machine. However, Virtual Machine movement can cause OmniVista to lose UDP traffic (e.g., SNMP Queries or Traps).

The VM Manager application is configured by clicking on one of the following tiles on the Home Page or links on the left side of the screen::

- **Hypervisor Systems** - Used to configure a Hypervisor (vCenter, XenServer, or Hyper-V).
- **VM Locator** – Enables the user to search and browse for VM Host Machines and Virtual Machines, and view Host Machine and Virtual Machine configurations.
- **Exclude VLAN** – Enables the user to configure a list of VM VLANs that VM Manager can ignore when polling Virtual Machine configurations.
- **VM VLAN Configuration** – Enables the user to configure VM VLANs and associate those VLANs with Universal Network Profiles (UNP) to monitor and manage Virtual Machines on the network.
- **VLAN Notification** - This panel displays a list of instances where Virtual Machines are mis-configured.
- **VMM Devices List** - Used to display information for switches connected to a Host Machine. Switches on this list are polled by VM Manager for VM Locator updates. Any switches connected to a Host Machine should be added to the VM Devices List to ensure the latest VM Locator information.
- **Settings** - Used to configure VM Polling and SBP configuration.

Virtualization/VM Manager Overview

The following sections provide an overview of virtualization and VM Manager's role in monitoring and managing Virtual Machines on the network. [Click here for VM Manager configuration steps.](#)

Virtualization

Virtualization is a way to use software (e.g., VMware) to create multiple Virtual Machines on a single Host Server to make better use of storage and server resources (CPU power, memory). These Virtual Machines run in isolation, side-by-side on the same physical machine. Each virtual machine can interact independently with other devices, applications, data and users as though it were a separate physical resource. This enables much more efficient and reliable use of server resources because different virtual machines can run different operating systems and multiple applications while sharing the resources of a single physical machine. And because each virtual machine is isolated from other virtualized machines, if one crashes, it does not affect the others. Moreover, Virtual Machines can dynamically migrate between Hosts to better utilize server resources.

Virtualization has a profound impact on the efficiency of the server farm. Without virtualization technology, the average server usage is about 10-20%, with virtualization, the average usage goes up to 40 - 60%. The introduction of virtualization in the data center has had a transforming effect on the way server farms are designed and operated. It also has implications for the network infrastructure.

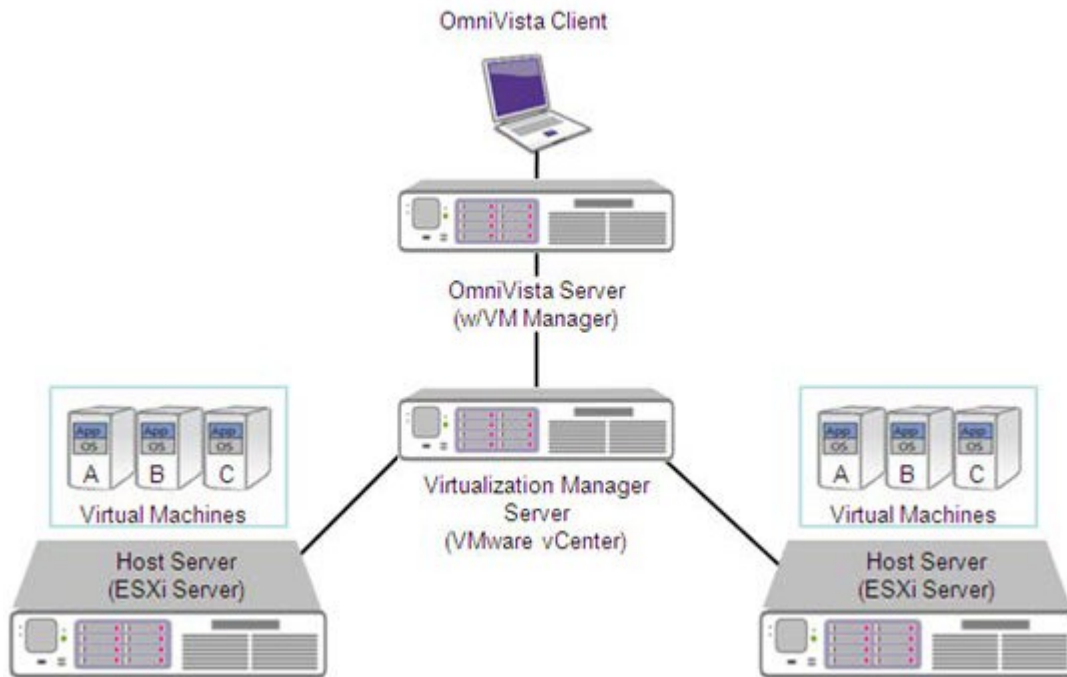
VM Manager and Virtualization

vCenter, XenServer, and Hyper-V provide a central location to monitor and manage Virtual Machines. OmniVista's VM Manager interacts with them to provide a unified view of virtual machines on the network. Although similar in operation, their configurations are slightly different. The sections below provide a high-level overview of their configuration and operation within the context of OmniVista's VM Manager. For detailed vCenter or XenServer configuration instructions, refer to the applicable vendor documentation.

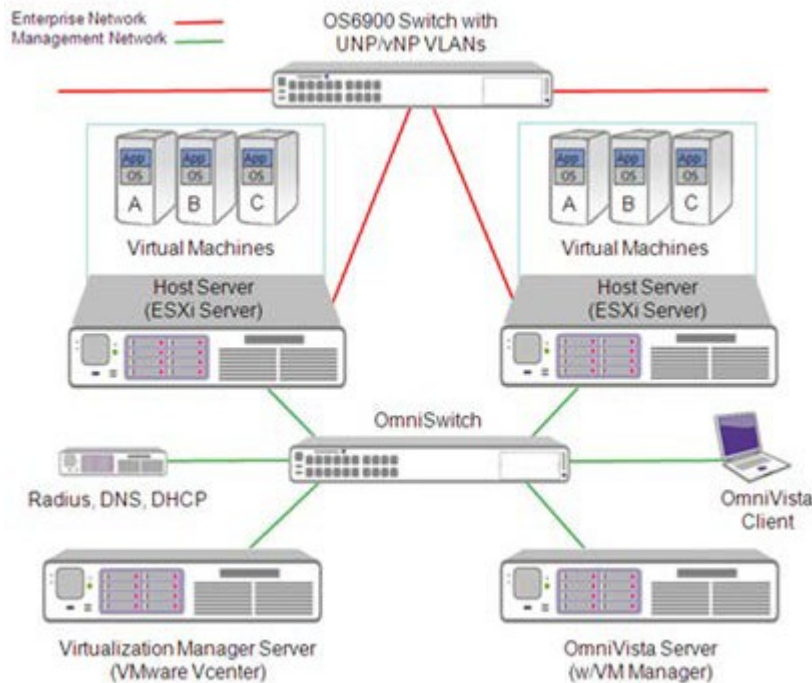
VMware vCenter

As shown below, the VM Manager application interfaces with VMware's vCenter to provide a unified view of Virtual Machines, their configurations, and designated switch configurations that, together enable proper traffic flows. Virtual Machines are configured on physical Host Servers that provide computing resources for the Virtual Machines. VMware's vCenter is a central service configured on its own physical server that interfaces with multiple Host Servers and the OmniVista Server through an OS6900 Switch. The diagram below shows a

configuration with a single vCenter. VM Manager will support up to two (2) vCenters.




The diagram below provides a high-level view of an OmniVista/vCenter network configuration.



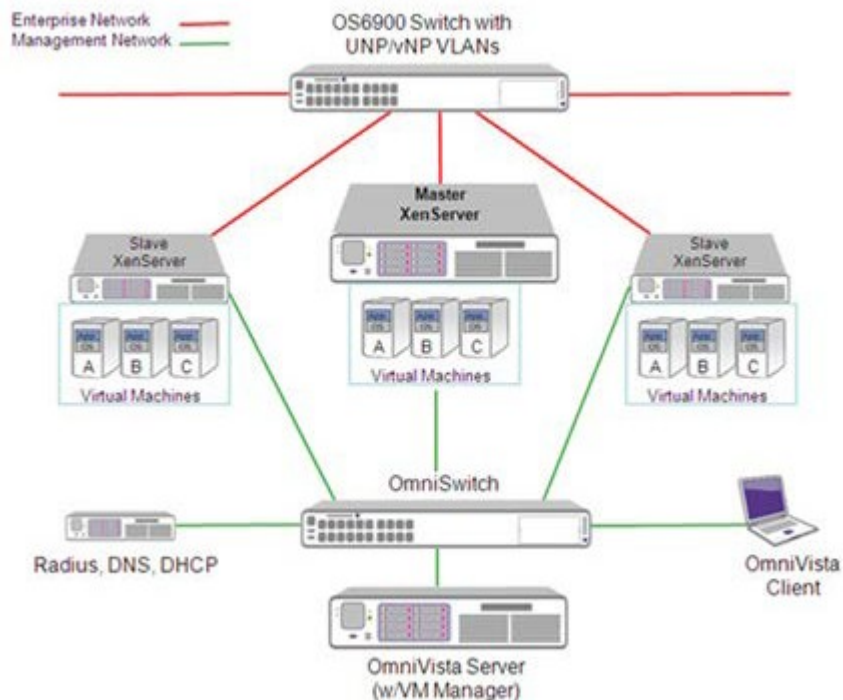
Citrix XenServer

The VM Manager application interfaces with XenServer to provide a unified view of Virtual Machines, their configurations, and designated switch configurations that, together enable proper traffic flows. Virtual Machines are configured on physical Host Servers that provide computing resources for the Virtual Machines. However, the XenServer is a central service that can be configured on any one of the Host Servers. OmniVista interfaces with a "Master" Host Server (Master XenServer), which provides the centralized view of the Virtual Machine configuration. The Master XenServer hosts Virtual Machines and can be connected to up to fifteen (15) "Slave" XenServers, which also host Virtual Machines. The Master XenServer then interfaces

with Slave Servers and the OmniVista Server through an OS6900 Switch. The diagram below shows a configuration with a single Master XenServer. VM Manager will support up to two (2) Master XenServers.



The diagram below provides a high-level view of an OmniVista/XenServer network configuration.



Microsoft Hyper-V

The VM Manager application interfaces with Hyper-V to provide a unified view of Virtual Machines, their configurations, and designated switch configurations that, together enable proper traffic flows. Virtual Machines are configured on physical Host Servers that provide computing resources for the Virtual Machines.

VM Manager and Traffic Shaping

VM Manager also interfaces with the Universal Network Profile (UNP) feature within OmniVista's Access Guardian application. The UNP Tag Rule feature enables you to assign VM VLANs to traffic shaping profiles based on UNP Classification Rules or Policy Lists. Any traffic matching the UNP Tag rule will have a UNP profile applied, and will be forwarded to a VLAN.

To utilize the VM Manager application, you first create Virtual Machine Port Groups inside the Host's networking configuration. The Virtual Machine Port Group is assigned a VLAN (VM VLAN) per your specification. Using UNP Tag rules on the switches, you can then associate VM VLANs with different UNPs (and their VLANs). Once a Virtual Machine is assigned a Virtual Machine Port Group, its network traffic is tagged with a VLAN number and the switch will know how handle the tagged packets based on the UNP Tag rule. This rule translates a VM VLAN to its corresponding UNP profile and VLAN on the switch.

If UNP Tag rules are consistently defined on all of the switches that carry VM network traffic, a Virtual Machine can move between Hosts connected to different switches without changes to its switch VLAN and traffic shaping parameters. The new switch will pick up the VM VLAN tag and know how to properly handle the VM network traffic. To help with consistency, the VM Notifications feature within VM Manager monitors Virtual Machine configurations and sends a notification in the event that a Virtual Machine configuration is missing.

Note: It is recommended that 'Management Network' that used to handle traffic for Host management, VM movement (vMotion), NFS, etc. be on a separate physical switch port and a separate UNP Tag rule or that the VLAN is defined differently from those defined for Virtual Machines. This requires a Host server with multiple NIC cards. This way, mis-configurations on a Virtual Machines' switch port will not cause any interruptions VMware's vCenter, Host Server, or OmniVista communications, and you can still use vSphere Client to manage Virtual Machines.

When Virtual Machines are created, you must also create a VLAN Tag for that machine. (Any traffic originating from that Virtual Machine will be tagged with that VLAN tag). To utilize the VM Manager application, you first create a UNP with a VLAN Tag Rule. Any traffic matching that VLAN tag is then routed to the VLAN associated with that UNP. You then create a VM VLAN and associate it with that UNP Profile and VLAN.

Configuring VM Manager

OmniVista's VM Manager application interfaces with a Hypervisor System (vCenter, XenServer, Hyper-V) to enable you to monitor Virtual Machines on your network. VM Manager also utilizes the Universal Network Profile (UNP) feature within OmniVista's Access Guardian application to apply UNP Rules to VM traffic. These UNP rules, which can be associated with QoS Policy Lists, are applied to UNP VLANs, and the traffic is then assigned to the applicable VLAN.

VM Manager Configuration Quick Steps

Configure Virtual Machines and vCenter as instructed by VMware. When you configure a Virtual Machine you must configure a VLAN Tag for the machine to enable VM Manager to monitor the machine and manage VM traffic. After configuring Virtual Machines and vCenter, follow the Quick Steps below to configure OmniVista's VM Manager. Generally, VM Manager configuration should use the following guidelines:

- Virtual Machines are tagged.
- Configure one UNP per VM VLAN.
- Configure a VLAN Tag Rule for each VM VLAN. (You can configure additional Classification Rules UNP or associate the UNP with Policy List, to further shape traffic.)

Note: VM Manager requires that the link discovery protocol be turned off on the port connecting the Hypervisor (ESXi Server/XenServer), or on the Hypervisor itself. Some Hypervisors may introduce LLDP packets which make it seem to have another physical bridging device, rather than an end station.

The recommended way to manage Virtual Machines in a data center using VM Manager is to have the Virtual Machines communicate using tagged VLAN packets, and provisioning the network using UNP VLAN Tag Classification rules over UNP Ports. Once all Virtual Machines are associated by VLAN tag with VM VLANs, any Virtual Machine movement will not require further adjustment to the configuration. This also ensures that OmniVista will notify the user through VM VLAN Notifications when a Virtual Machine and its VM VLAN are mis-configured.

1. In the VM Manager application, go to the Hypervisor Systems Screen to configure VM Manager's connection to a Hypervisor (vCenter, XenServer, Hyper-V).

Important Note: The system time on all hypervisors must be correctly set and synchronized with the time on the OmniVista Server.

2. Go the Unified Access application and create a VXLAN Mapping Template (Unified Access - Unified Profile - Template - VXLAN Mapping).

3. In the Unified Access application, create an Access Classification Rule for the VLAN (Unified Access - Unified Profile - Template - Access Classification).

4. Assign the UNP Policy to network switches.

Hypervisor Systems

The VM Manager Hypervisor Systems Screen displays a list of all VM Servers connected to OmniVista; and is used to configure the connection from OmniVista to a VM Server. You can also edit and delete VM Servers. OmniVista supports VMware's vCenter, Citrix XenServer, and Microsoft Hyper-V.

If no VM Servers have been configured, the Hypervisor Systems Table is not displayed. Click on a link (VMware's vCenter, Citrix XenServer, Microsoft Hyper-V) at the top of the screen to bring up the Create Hypervisor Systems Screen to create a VM Server.

Important Note: The system time on all hypervisors must be correctly set and synchronized with the time on the OmniVista Server.

Creating a VM Server

To configure the connection from OmniVista to a VM Server, click on the Create icon **+** and complete the fields as described below. The fields differ slightly depending on the server type selected (VMware's vCenter, Citrix XenServer, Microsoft Hyper-V). When you have finished completing the fields, click on the **Create** button.

Note: You can test the connection to the VM Server before creating it by clicking on the **Test Connection** button after completing the fields.

VMware vCenter

- **VM Server Type** - Select **VMware vCenter** from the drop-down menu).
- **URL** - The IP address of the VM Server. For a vCenter Server, enter the IP address, followed by "/sdk" (e.g., https://10.255.11.1/sdk).
- **Name** - User-configured name for the VM Server.
- **User** - Administrator's User Name.
- **Password** - The password needed to access the VM Server.
- **Re-Type Password** - Re-Type password needed to access the VM Server.


Citrix XenServer

- **VM Server Type** - Select **Citrix XenServer** from the drop-down menu).
- **URL** - The IP address of the VM Server. You must add (e.g., https://10.255.11.1).
- **Name** - User-configured name for the VM Server.
- **User** - Administrator's User Name.
- **Password** - The password needed to access the VM Server.
- **Re-Type Password** - Re-Type password needed to access the VM Server.

Microsoft Hyper-V

- **VM Server Type** - Select **Microsoft Hyper-V** from the drop-down menu). The server type you selected on the previous screen will automatically be selected. However, you can also select a different server type from this drop-down menu.
- **IP Address** - The IP address of the VM Server (e.g., https://10.255.11.1).
- **Name** - User-configured name for the VM Server.
Domain - The Hyper-V Domain.
- **User** - Administrator's User Name.
- **Password** - The password needed to access the VM Server.
- **Re-Type Password** - Re-Type password needed to access the VM Server.

Editing a VM Server

Select a VM Server in the table and click on the Edit icon . Edit the field(s) as described above and click on the **Apply** button. You can only edit the server password.

Deleting a VM Server

Select a VM Server(s) in the table, click on the Delete icon  then click **OK** at the Confirmation Prompt.

Hypervisor Systems Table

The Hypervisor Systems Table displays a list of all VM Servers connected to OmniVista.

- **Name** - User-configured name for the VM Server.
- **URL** - The IP address of the VM Server.
- **VM Server Type** - The VM Server type (VMware's vCenter, Citrix XenServer, and Microsoft Hyper-V).
- **User** - Administrator's User Name.
- **Status** - The administrative status of the Server (Up, Down, Unknown).

VM Locator - Host Networks

The VM Manager Host Networks Screen is used to search for and display information on the Host Machines on which the Virtual Machines reside. You can view information for all Host Machines or enter search criteria to view specific Host Machines. The information displayed is based on the most recent search. To refresh the information with the most recent data, repeat the search.

Searching for Host Machines

Select an option from the **Search By** drop-down menu (e.g., Host MAC Address, Host IP Address) and enter the search criteria. (You can also just select "All Hosts" to search for all Host Machines.) By default, OmniVista will conduct an historical search. If you want to do a live search, set the **Live Search** slider to **Enabled**. You can also enable "Stop after 1st Match" to stop the search after finding the first match. When you are done, click on the **Apply** button. The results are displayed in the Host Networks Table.

Note: If you enable the "Stop after 1st Match" option, OmniVista will stop searching after at least one match is found; however more than one match may be displayed.

Host Network Table

The Host Networks Table displays basic information about all configured VM Server networks. To view detailed information, click on an entry in the table.

Basic Information

- **Hypervisor Host** - The user-configured name for the Host Machine. If none is configured, the IP address of the Host is displayed.
- **VM Server** - The user-configure name of the VM Server (vCenter, XenServer, Hyper-V).
- **IP Address** - The Host Machine IP address.
- **Network Mask** - The corresponding network mask of the Host Machine.
- **Network Name** - The VM VLAN that the Virtual Machine's network interface is associated with. The network traffic for a Port Group may be tagged or untagged.
- **Number of Networks** - The number of networks (also known as Physical Interfaces (PIFs) on the Host Machine.
- **Switch IP Address** - The IP address of the switch to which the Host Machine is connected.
- **Slot/Port** - The slot/port of the switch to which the Host Machine is connected.
- **Port VLAN** - The VLAN or SPB Service ID that the switch uses to classify Virtual Machine network traffic.
- **UNP** - The UNP associated with the Host Machine's interface.
- **Locator Time** - The time the Virtual Machine's network traffic was detected on the switch port.
- **Last Update** - The date and time the Host Machine configuration was last updated.

Detailed Information

- **Hypervisor Host** - The user-configured name for the Host Machine. If none is configured, the IP address of the Host is displayed.
- **VM Server** - The user-configure name of the VM Server (vCenter, XenServer, Hyper-V).
- **IP Address** - The Host Machine IP address.
- **Network Mask** - The corresponding network mask of the Host Machine.
- **Network Name** - The VM VLAN that the Virtual Machine's network interface is associated with. The network traffic for a Port Group may be tagged or untagged.

- **Number of Networks** - The number of networks (also known as Physical Interfaces (PIFs) on the Host Machine.
- **Switch IP Address** - The IP address of the switch to which the Host Machine is connected.
- **Slot/Port** - The slot/port of the switch to which the Host Machine is connected.
- **Port VLAN** - The VLAN or SPB Service ID that the switch uses to classify Virtual Machine network traffic.
- **UNP** - The UNP associated with the Host Machine 's interface.
- **Locator Time** - The time the Virtual Machine's network traffic was detected on the switch port.
- **Last Update** - The date and time the Host Machine configuration was last updated.
- **Uptime** - The amount of time the Host Machine has been up (time since last reboot).
- **CPU Count** - The number of processors on the Host Machine.
- **CPU Model** - The model name of the Host Machine CPU.
- **Service ID** - The Service ID that the switch uses to classify Virtual Machine network traffic.
- **ISID** - The ISID that the switch uses to classify Virtual Machine network traffic.
- **Port Status** - The operational status of the Host port connected to the switch.
- **Port Speed** - The speed of the Host port connected to the switch.
- **Duplex** - The duplex mode (half duplex, full duplex, or auto duplex) of the Host port connected to the switch.
- **Disposition** - The switch port's disposition (Bridging/Filtering).
- **Classification Source** - The Classification Policy under which the device was learned.
- **Data Center** - The name of the Data Center to which the Virtual Machine is assigned.
- **Cluster** - The Cluster in which the Virtual Machine resides.
- **Memory Usage** - The amount of RAM currently being used by the Host Machine, in MB.
- **Vendor** - The Manufacturer of the Host Machine (e.g., HP).
- **Status** - The administrative status of the Host Machine.
- **VM Motion Enabled** - Whether or not VM Motion is enabled. If enabled, Virtual Machines can be moved from one Host to another and the VM configuration will be dynamically updated.
- **CPU** - The speed of the Host Machine CPU, in MHz.
- **Memory Size** - The amount of RAM on the Host Machine, in MB.
- **Num of CPU Pkgs** - Number of physical CPU packages on the Host Machine.
- **Num of Threads** - The number of threads.
- **Num of HBAs** - The number of Host Bus Adapters (HBA).
- **CPU Usage** - The amount of CPU currently being used by the Host Machine, in MHz.
- **Power** - The Power status of the Host Machine (Powered On, Powered Off, Suspended).
- **DNS Name** - The name of the DNS associated with the Host Machine (if applicable).

VM Locator - VM Networks

The VM Manager VM Networks Screen is used to search for and display information on the Virtual Machines residing on the Host Machine. You can view information for all Virtual Machines or enter search criteria to view specific Virtual Machines. The information displayed is based on the most recent search. To refresh the information with the most recent data, repeat the search.

Searching for Virtual Machines

Select an option from the **Search By** drop-down menu (e.g., VM MAC Address, VM IP Address) and enter the search criteria. (You can also just select "All VMs" to search for all Virtual Machines.) By default, OmniVista will conduct an historical search. If you want to do a live search, set the **Live Search** slider to **Enabled**. You can also enable Stop after 1st Match to stop the search after finding the first match. When you are done, click on the **Apply** button. The results are displayed in the VM Networks Table.

Note: "Live Search" locates the most recent location of VMs by doing live query against known switches. If a location is found for a VM's MAC address, information is displayed with the latest timestamp. If a live search does not result in a match, the last historical location and past timestamp is displayed. Historical location will produce multiple results from the past for tracing purposes. Historical search results can also display "false positive" information. For example, a non-uplink port used to connect an end station that has become an uplink port can still be displayed. Furthermore, uplink status that may be learned later will not be used to determine location. This condition will correct itself but historical data will be persistent regardless. Use the Locator timestamp to determine if the information is noteworthy.

Note: If you enable the "Stop after 1st Match" option, OmniVista will stop searching after at least one match is found; however more than one match may be displayed.

VM Networks Table

The VM Networks Table displays basic information about all configured Virtual Machines. To view detailed information, click on an entry in the table. As indicated below, the information varies slightly depending on server type (vCenter, XenServer, Hyper-V).

Basic Information

- **VM Name** - The user-configure Virtual Machine name.
- **DNS Name** - The name of the DNS associated with the Virtual Machine (if applicable).
- **MAC Address** - The MAC address of the Virtual Machine.
- **IP Address** - The IP address of the Virtual Machine.
- **Network Name** - The VM VLAN that the Virtual Machine's network interface is associated with. The network traffic for a Port Group may be tagged or untagged.
- **VM Server** - The user-configure name for the VM Server.
- **Hypervisor Host** - The IP address of the Host Machine on which the Virtual Machine resides.
- **Switch IP Address** - The IP address of the switch to which the Host Machine is connected.
- **Slot/Port** - The slot/port of the switch to which the Host Machine is connected.
- **VM VLAN** - The Tag Value for the VM VLAN in the Host System.
- **UNP** - The UNP associated with the Virtual Machine.
- **Locator Time** - The time the Virtual Machine's network traffic was detected on the switch port.
- **Last Update** - The date and time the Host Machine configuration was last updated.

Detailed Information


- **VM Name** - The user-configure Virtual Machine name.
- **DNS Name** - The name of the DNS associated with the Virtual Machine (if applicable).
- **MAC Address** - The MAC address of the Virtual Machine.
- **IP Address** - The IP address of the Virtual Machine.
- **Network Name** - The VM VLAN that the Virtual Machine's network interface is associated with. The network traffic for a Port Group may be tagged or untagged.
- **VM Server** - The user-configure name for the VM Server.
- **Hypervisor Host** - The IP address of the Host Machine on which the Virtual Machine resides.
- **Switch IP Address** - The IP address of the switch to which the Host Machine is connected.
- **Slot/Port** - The slot/port of the switch to which the Host Machine is connected.
- **VM VLAN** - The Tag Value for the VM VLAN in the Host System.

- **UNP** - The UNP associated with the Virtual Machine.
- **Locator Time** - The time the Virtual Machine's network traffic was detected on the switch port.
- **Last Update** - The date and time the Host Machine configuration was last updated.
- **Address Type** - The Virtual Machine address type (Assigned/Unassigned).
- **Guest OS** - The operating system of the Virtual Machine.
- **Power** - The Power status of the Virtual Machine (Powered On, Powered Off, Suspended).
- **Up Time** - The amount of time the Virtual Machine has been up (time since last reboot).
- **Port VLAN** - The VLAN that the switch uses to classify Virtual Machine network traffic.
- **Service ID** - The SPB Service ID that the switch uses to classify Virtual Machine network traffic.
- **ISID** - The ISID that the switch uses to classify Virtual Machine network traffic.
- **Port Status** - The operational status of the Virtual Machine port connected to the Host Machine.
- **Port Speed** - The speed of the Virtual Machine port connected to the Host Machine.
- **Duplex** - The duplex mode (half duplex, full duplex, or auto duplex) of the Virtual Machine port connected to the Host Machine.
- **Disposition** - The switch port's disposition (Bridging/Filtering).
- **Classification Source** - The Classification Policy under which the device was learned.
- **Status** - The operating status of the Virtual Machine.
- **Network Usage** - The percentage of network resource being used by the Virtual Machine.

Exclude VLAN


The VM Manager Exclude VLAN Screen displays a list of all Exclude VM VLANs. It is also used to create, edit and delete Exclude VLANs. When OmniVista polls a VM Server, it checks the Virtual Machine configuration and sends a notification to VM Manger if there is a problem with the configuration (displayed on the VLAN Notifications Screen). The Exclude VLAN Screen is used to define VM VLANs that should be ignored by OmniVista when conducting VM polling (e.g., the VM Network Management VLAN). VLANs listed here will be ignored during support checks.

Creating an Exclude VLAN


Create icon  and complete the fields as described below.

- **VLAN ID** - Enter a VLAN ID, multiple VLAN IDs or a range of VLANs.
- **Description** - Enter a description for the Exclude VLAN(s).

Editing and Exclude VLAN

Select a VLAN in the Exclude VLAN List and click on the Edit icon . You can only edit the "Description" field. Edit the field and click on the **Apply** button.

Deleting an Exclude VLAN

Select a VLAN(s) in the Exclude VLAN List, click on the Delete icon , the click on **OK** at the Confirmation prompt.

Exclude VLAN List

The Exclude VLAN List displays all configured Exclude VLANs.

- **VLAN ID** - The VLAN ID.
- **Description** - User configured description for the VLAN.

VM VLAN Configuration

The VM Manager VM VLAN Configuration Screens are used to associate VM VLANs with Universal Network Profiles (UNP), and enable MVRP Ports and One-Touch SPB on network switches/ports.

The recommended way to manage Virtual Machines in a data center using VM Manager is to have the Virtual Machines communicate using tagged VLAN packets, and provisioning the network using UNP VLAN Tag Classification rules over UNP Ports. Once all Virtual Machines are associated by VLAN tag with VM VLANs, any Virtual Machine movement will not require further adjustment to the configuration. This also ensures that OmniVista will notify the user through VM VLAN Notifications when a Virtual Machine and its VM VLAN are mis-configured.

Note: SPB is supported on OS10K and OS6900 Switches running AOS 7.3.1.R01 and later, with an *Advanced* License. To support a mixture of devices using the same screen, OmniVista only pushes configurations which are applicable for specific device types, skipping the rest. For SPB-capable devices, all attributes are applicable. For non-SPB devices, SPB-specific attributes will be skipped and only regular bridging UNP changes will be updated.

VM VLAN Configuration - Apply UNP VLAN

The VM Manager Apply UNP VLAN Screen is used to apply VLAN Tag Rules to network switches/ports. The recommended way to manage Virtual Machines in a data center using VM Manager is to have the Virtual Machines communicate using tagged VLAN packets, and provisioning the network using UNP VLAN Tag Classification rules over UNP Ports. Once all Virtual Machines are associated by VLAN tag with VM VLANs, any Virtual Machine movement will not require further adjustment to the configuration. This also ensures that OmniVista will notify the user through VM VLAN Notifications when a Virtual Machine and its VM VLAN are mis-configured.

Complete the fields as described below and click on the **Apply** button to apply VLAN Tag Rules to network switches/ports.

- **Select VLAN Tag Rules** - Select an existing VLAN Tag Rule(s) from the drop-down list. You can also click on the Add icon **+** to go to the Access Guardian application and create a new rule(s).
- **Enable UNP Ports** - Enables/Disables UNP on the selected ports. By default, UNP Ports are enabled when you create UNP Port Policies.
- **Select UNP Port Policy** - Select as UNP Port Policy from the drop-down menu. You can also click on the Add icon **+** to go to the Access Guardian application and create a new policy. You can only select one policy.
- **Select Devices** - Select an option from the drop-down menu "Use Switch Picker" or "Use Topology" and click on the **Add/Remove Devices** button to select the device(s) on which you want to apply the VLAN Tag Rule(s).
- **Select Ports** - Click on a device and click on the **Add/Remove Ports** button to select the port(s) on which you want to apply the VLAN Tag Rule(s).

VM VLAN Configuration - Enable MVRP Ports

The VM Manager Enable MVRP Ports Screen is used to enable MVRP Ports on network switches. Select an option from the drop-down menu "Use Switch Picker" or "Use Topology" and click on the **Add/Remove Devices** button to select the device(s) on which you want to enable MVRP Ports. Then select a device and click on the **Add/Remove Ports** button to enable MVRP on those ports.

VM VLAN Configuration - Enable One-Touch SPB

The VM Manager Enable One-Touch SPB Screen is used to enable SPB Interfaces on network switches. One-Touch SPB configures switches for backbone communication to transport VM traffic that is classified to service domain using UNP for SPB. UNP for SPB specifies which Service Access Port and which Policy List will be used. Subsequently, it will determine the correct layer 2 domain for VM data.

Select an option from the drop-down menu "Use Switch Picker" or "Use Topology" and click on the **Add/Remove Devices** button to select the device(s) on which you want to enable MVRP Ports. Then select a device and click on the **Add/Remove Interfaces** button to enable SPB on those interfaces.

VLAN Notification

The VM Manager VLAN Notification Screen displays VM VLAN Notifications generated by the VMM Service for missing VLAN/UNP configuration on a switch slot/port where VMs are connected. The notifications briefly describe the problem and enable you to resolve it. Ideally, you want to have no notifications in this panel. You can also resolve configuration problems using the Resolve Feature. There are two links on the screen:

- **Active Notifications** - Displays all active notifications. Note that a user can move a notification to the Ignored Notifications List selecting the notification(s) in the Active Notifications List and clicking on the **Ignore** Button.
- **Ignored Notifications** - Displays any notifications that the user has moved to the tab because they may describe alternate configurations that are known by the user to work. Generally, a user will move notifications to the Ignored List because they have alternate configurations (e.g., MAC, MAC Range, IP) which are known to work, in addition to the supported UNP VLAN Tag configurations. Note that a user can also move the Ignored notifications back to the Active list by highlighting the notification(s) on the Ignored Tab and clicking on the **Activate** Button.

You can also select the notification and click on the **Resolve** button to open the UNP VLAN Wizard and resolve the problem.

Resolving a VM Configuration Problem

When a mis-configuration notification appears in the VLAN Notification window, select the notification and click the Resolve button to bring up the UNP VLAN Configuration Wizard. Use the Wizard to correct the configuration problem described in the notification (e.g., missing tag rule, slot/port).

If a VM discovered on a UNP or SPB access port and OmniVista's VLAN tag rule has not been created on the switch(es), OmniVista create an entry in the Active Notifications List. The "Resolve" button will guide you through the necessary configuration steps to resolve the problem and clear the Notification.

For SPB configurations, OmniVista may generate a new profile with an ISID derived from the VM VLAN and associate this value into one of the 4 pre-configured SBP BVLANS, if it has not been created. This profile is used to classify the VLAN's traffic into a the provider backbone's BVLAN. OmniVista uses a round-robin mechanism to associate BVLAN to a VM VLAN and pushes these parameters into an SPB profile. A VM VLAN will have a uniquely associated ISID created by adding the VM VLAN ID to the starting ISID number defined in SPB Settings Screen.

If a VLAN Tag rule already exists as it may pertain to regular bridging UNP, OmniVista will append an additional SPB Profile name using the corresponding SPB Profile, which was auto-generated for the VLAN. With that in mind, "ag Position" will be specified as "Outer" and the SPB Profile will be the one generated based on the VM VLAN. The Customer Domain ID will be based on the ports' Customer Domain ID. For example, If VLAN 31 is detected on SPB ports with Customer Domains of 1 and 2, 2 entries of UNP VLAN will

be created - VLAN 31, Customer Domain 1> and VLAN 31, Customer Domain 2 - and corresponding SPB Profiles will be created in Access Guardian.

UNP Profiles with blank policy lists will be created if necessary to resolve a UNP VLAN tag rule and the rule will be specified in the UNP VLAN List. If the UNP Profile already exists, it will be automatically picked up in creating UNP VLAN List entry. Auto-generated UNP Profiles have the name UNP XX where XX is the VLAN ID. The Default UNP Profile is named UNP1 and its VLAN ID is 1. This UNP will be used to resolve "Default UNP" notifications for regular UNP ports.

Note: In both cases where bridging-domain UNP Profiles or SPB Profiles are generated, the Policy List is left blank. Having these profiles will only facilitate connectivity. You can later decide to modify the UNP Profile to that rule by clicking on Edit button and reassigning the policy list.

Note: SPB Profiles that are auto-generated using the Resolve function require that "One Touch SPB" is configured successfully on the switch(es) and the same default configuration parameters exist that existed when "One Touch SPB" was first configured. If this configuration has changed, "One Touch SPB" needs to be re-invoked.

Notification List

The following fields are displayed on both the Active and Ignored Notifications Lists.

- **VLAN** - The VM VLAN ID.
- **Host Name** - The name of the Host Machine hosting the Virtual Machine.
- **Switch** - The switch connected to the Host Machine.
- **Slot/Port** - The slot/port number of the port connecting the Host Machine to the switch.
- **Missing Configuration** - A brief description of the configuration problem.
- **Missing Configuration Slot/Port** - The port(s) missing from the configuration.
- **Port Groups** - The port groups missing from the configuration.
- **Last Update** - The time the notification was last updated.
- **Create Time** - The time the notification was created.

VMM Devices List

The VM Manager VMM Devices List Screen displays information for switches connected to a Host Machine. Switches on this list are polled by VM Manager for VM Locator updates. Any switches connected to a Host Machine should be added to the VM Devices List to ensure the latest VM Locator information. You can add/remove switches from the list by clicking on the **Select Devices** button at the top of the screen and selecting devices for the list.

VMM Devices List

- **Friendly Name** - User-configured name for the device.
- **Name** - The name of the device.
- **Address** - The address of the device.
- **Status** - The operational status of the device. It displays "Up" if the device is up and responding to polls. It displays "Down" if the device is down and not responding to polls. It displays "Warning" if the switch has sent at least one warning or critical trap and is thus in the warning state.
- **DNS Name** - The DNS name of the device.
- **Type** - The type of device chassis (e.g., OS6850-24).

- **Version** - The version number of the device software (e.g., 6.6.5.96.R02). OmniVista may not be able to determine the software version on some third-party devices. In these cases, the field will be blank.
- **Location** - The physical location of the device (e.g., Test Lab).
- **NOD** - Whether or not the device is an NOD device (Yes/No).
- **Activated Licenses** - Activated optional licenses on the device.
- **FTP User Name** - The CLI/FTP user name for the device.
- **SNMP Version** - The SNMP version that OmniVista uses to communicate with the device.
- **v1/2 Read Community** - The device's SNMP v1/2 "get" community name, if applicable.
- **v1/2 Write Community** - The device's SNMP v1/2 "set" community name, if applicable.
- **v3 User Name** - The device's SNMP v3 user name, if applicable.
- **Last Upgrade Status** - The status of the last firmware upgrade on the device.
 - "Successful" - Successful BMF and Image upgrade performed.
 - "Successful (BMF)" - Successful BMF upgrade performed.
 - "Successful (Image)" - Successful Image upgrade is performed.
 - "Failed (BMF, Image)" - BMF and Image upgrade failed.
 - "Failed (BMF)" - BMF upgrade failed.
 - "Failed (Image)" - Image upgrade failed.
- **Backup Date** - The date that the device's configuration and/or image files were last backed-up to the OmniVista Server.
- **Backup Version** - The firmware version of the configuration and/or image files that were last backed-up to the OmniVista Server.
- **Last Known Up At** - The date and time when the last poll was initiated on the device.
- **Description** - A description of the device, usually the vendor name and model.
- **Traps** - The status of trap configuration for the device. "On" means that traps are enabled. "Off" means that traps are disabled. "Not Configurable" means that traps for this device are not configurable from OmniVista. (Note that traps may have been configured for such devices outside of OmniVista.) "Unknown" means that OmniVista does not know the status of trap configuration on this device.
- **Seen By** - The User Groups that are able to view the device. OmniVista is shipped with the following pre-defined user groups (Default, Writers, Network Administrators, Administrators) that have different security permissions.
- **Running From** - For AOS devices, this field indicates whether the switch is running from the Certified directory or from the Working directory. This field is blank for all other devices. For AOS devices, the directory structure that stores the switch's image and configuration files in flash memory is divided into two parts:
 - The Certified directory contains files that have been certified by an authorized user as the default configuration files for the switch. When the switch reboots, it will automatically load its configuration files from the certified directory if the switch detects a difference between the certified directory and the working directory.
 - The Working directory contains files that may or may not have been altered from those in the certified directory. The working directory is a holding place for new files to be tested before committing the files to the certified directory. You can save configuration changes to the working directory. You cannot save configuration changes directly to the certified directory.

Note that the files in the certified directory and in the working directory may be different from the running configuration of the switch, which is contained in RAM. The running configuration is the current operating parameters of the switch, which are originally loaded from the certified or working directory but may have been modified through CLI commands, WebView commands, or OmniVista.

Modifications made to the running configuration must be saved to the working directory (or lost). The working directory can then be copied to the certified directory if and when desired.

Note: OmniVista supports the Multiple Working Directories Feature available on OS10K and OS6900 Switches (AOS Release 7.2.1.R01 and later). This feature allows the user to create multiple Working Directories on the switch that can be used to save specific switch configurations. The user can create any name for these "Working" Directories (e.g., "Marketing Switch 05-23-15"). If the switch is running from one of these user-created directories, the directory name is displayed in this field.

- **Changes** - For AOS devices, this field indicates the state of changes made to the switch's configuration. This field is blank for all other devices. This field can display the following values:
 - **Certified** - Changes have been saved to the working directory, and but the working directory has been copied to the certified directory. The working directory and the certified directory are thus identical.
 - **Uncertified** - Changes have been saved to the working directory, but the working directory has not been copied to the certified directory. The working directory and the certified directory are thus different.
 - **Unsaved** - Changes have been made to the running configuration of the switch that have not been saved to the working directory.
 - **Blank** - When this field is blank for an AOS device, the implication is that OmniVista knows of no unsaved configuration changes and assumes that the working and certified directories in flash memory are identical.
- **Discovered** - The date and time when OmniVista successfully pings or polls the switch for the first time. This value remains unchanged until the switch entry is deleted. This field will remain blank if OmniVista does not ping or poll the switch at all.
- **No. of Licenses Used** - The total number of Core (AOS) or Third-Party licenses being used. For example, a stack of 4 switches would require 4 licenses, a VC of 6 would require 6 licenses. If a stack splits, the number of licenses reserved for the device before the split is maintained even though modules have been reduced to less than 5. This way, the license counts are reserved for the stack to recover.
- **License Type** - The type of license used by the device (e.g., AOS, Third Party).

Settings - VM Polling

The VM Manager VM Polling Screen is used to set the interval at which OmniVista will poll VM Servers. The interval you set should be determined based on the number of Virtual Machines you are managing. The more machines you are managing, the more resource-intensive the operation will be. The Polling Interval should generally be the same as the interval set for "Regular Updates" in the Setting Frequencies Screen in the Discovery Application.

Select an option from the drop-down menu (Minutes, Hours, Days), enter a Polling Interval and click on the **Apply** button. You can also click on the **Poll Now** button to perform an immediate poll of all connected VM Servers.

Settings - SBP

The VM Manager Settings Screen enables you to modify the default "One Touch SPB" configuration. You can overwrite the system-chosen starting ISID and create a value unique to an OmniVista Server. This enables you to create different L2 tunneling domains for Virtual Machines (VMs). VMs associated with the same VLAN/Network in the Hypervisor's environment but placed on different ISIDs will not communicate with each other. Other attributes can also be customized, including Control BVLAN, BVLANS 2 - 4, and ECT ID 1 - 4, which are tie-breaking algorithms. Update any fields as described below and click on the **Apply** button.

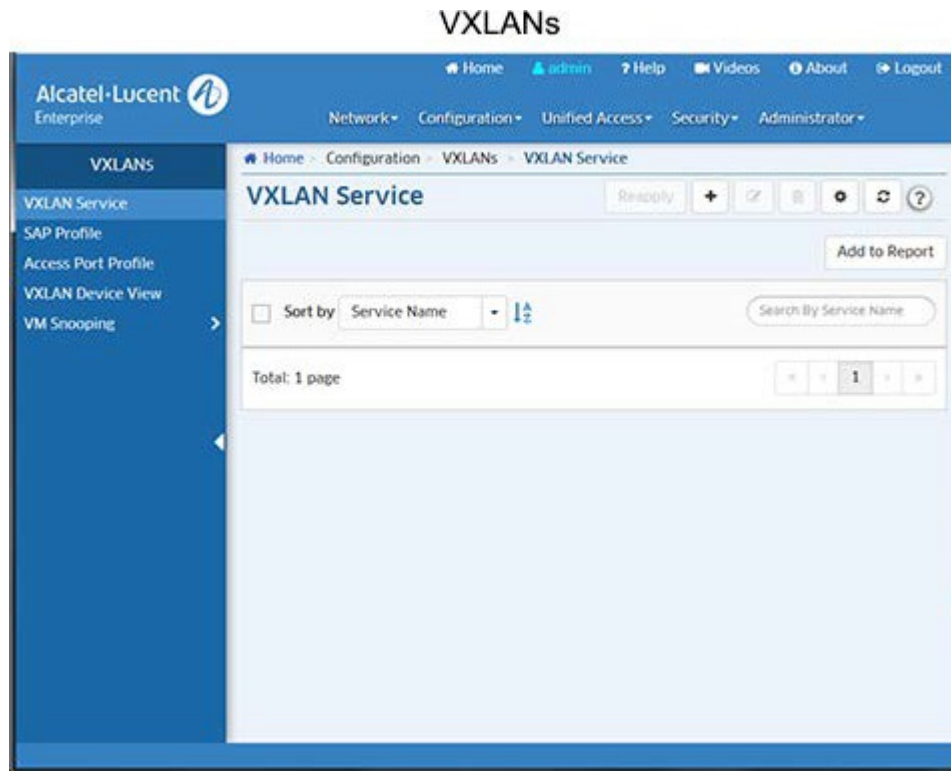
- **Starting ISID** - Used in the auto-generation of SPB Profiles to resolve any notification that is raised when the user clicks on the "Resolve" button in the VM Manager VLAN Notification node. To keep consistent mappings of VLAN-to-ISID, the VM VLAN number is added to the starting ISID number to determine the ISID for each VLAN.
- **Control BVLAN** - One of the four (4) BVLANS created in "One-Touch SPB" to take advantage of shortest path bridging topology from the source. Used for traffic as well as control information.
- **ECT ID 1** - Cost Tree Identifier (ECT ID) assigned to Control BVLAN. The ECT ID assigns a tie-breaking algorithm to the BVLAN that is used for Shortest Path Tree (SPT) calculations
- **Additional BVLAN 2** - One of the four (4) BVLANS created in "One-Touch SPB" to take advantage of shortest path bridging topology from the source. Used only for traffic.
- **ECT ID 2** - Cost Tree Identifier (ECT ID) assigned to BVLAN 2. The ECT ID assigns a tie-breaking algorithm to the BVLAN that is used for Shortest Path Tree (SPT) calculations
- **Additional BVLAN 3** - One of the four (4) BVLANS created in "One-Touch SPB" to take advantage of shortest path bridging topology from the source.
- **ECT ID 3** - Cost Tree Identifier (ECT ID) assigned to BVLAN 3. The ECT ID assigns a tie-breaking algorithm to the BVLAN that is used for Shortest Path Tree (SPT) calculations
- **Additional BVLAN 4** - One of the four (4) BVLANS created in "One-Touch SPB" to take advantage of shortest path bridging topology from the source.
- **ECT ID 4** - Cost Tree Identifier (ECT ID) assigned to BVLAN 4. The ECT ID assigns a tie-breaking algorithm to the BVLAN that is used for Shortest Path Tree (SPT) calculations.

28.0 VXLANs

Virtual Extensible LAN (VXLAN) is a network virtualization scheme tailored to address the evolving trends such as server virtualization and cloud computing in the current data center deployments. Legacy Layer 2 (L2) bridging using VLANs to segregate user traffic is no longer sufficient to meet the scale of current and future requirements. VXLAN uses a L2 over L3 encapsulation technique to overlay L2 network segments on L3 network infrastructure. Basically, a MAC frame received is encapsulated in an IP packet with a UDP header and sent out over the L3 network. Use of L3 as the transport network automatically allows use of ECMP routes, thus increasing utilization of the network infrastructure that cannot be matched by bridged networks using STP.

VXLAN supports up to 16 million virtualized L2 segments called VXLAN segments. Only servers, or more precisely VMs residing in the servers, that are attached to the same VXLAN segment can communicate with each other. Each VXLAN segment is identified by a 24 bit segment ID called VXLAN Network Identifier (VNI). The VNI identifies the scope of the inner MAC frame originated by the individual VM. Thus, we can have overlapping MAC addresses across segments but never have traffic "crossover" since the traffic is isolated using the VNI. The VNI is in an outer header which encapsulates the inner MAC frame originated by the VM. The VMs themselves are completely unaware of the VXLAN and VNI and communicate with each other as if connected directly over a VLAN-based L2 network.

Note: VXLANs are only supported on OS6900-Q32 and OS6900-X72 Switches.



VXLAN Service

The VXLAN Service Screen displays all configured VXLAN Services, and is used to create, edit, and delete VXLAN Services. A VXLAN Service defines a Virtual Forwarding Instance (VFI) that is capable of learning device MAC addresses from the access side and from the network side and then switching the traffic based on this information. Each VXLAN Service is basically an VFI that is capable of learning customer MAC addresses from the access side (Service Access Points - SAP) and from the network side (mesh Service Distribution Point - SDP) and then switching traffic based on this information.

Creating a VXLAN Service

Click on the Create icon **+** and complete the fields as described below to create the VXLAN Service.

- **Service Name** - The name of the VXLAN Service (up to 32 characters).
- **VNID** - The Virtual Network Identifier (VNID) is a 24-bit segment ID (also referred to as a VXLAN Segment ID) that is used to identify encapsulated frames. A VNID is bound to a VXLAN Service when the service is created. OmniVista will auto-generate a unique VNID for VXLAN Service if VNID is set to zero (0).
- **VXLAN Network Profile** - The VXLAN Network Profile associated with the VXLAN Service. The profile specifies the UDP Port and VRF Name. If necessary, create a VXLAN Network Profile.
- **VLAN Translation** - Enables/Disables egress VLAN translation for all SAPs associated with the VXLAN Service. Enabling translation at the service level is only applicable if the corresponding access ports for the SAPs also have VLAN translation enabled.
- **Admin Status** - Enables/Disables the administrative status of the VXLAN Service. Disable the administrative status of the service and any associated SAPs and SDPs before deleting a service.

Note: The **Reapply** button is enabled and can be used (in certain cases) when a VXLAN Service create/edit fails. Click here for more information.

Creating a VXLAN Network Profile

Click on the Create icon **+** and complete the fields as described below to create the VXLAN Network Profile.

- **Profile Name** - The VXLAN Network Profile Name. Select a profile from the drop-down menu, or click on the Add icon **+** to go to the Screen and create a profile. A “Default” profile with default parameters is available. This profile cannot be modified or deleted; however if a VXLAN Service is deleted, this profile will be deleted if no other service is using it.
- **VRF Name** - The VRF instance associated with the profile.
- **UDP Port** - The UDP Port used by the VXLAN Service.

Creating an Service Distribution Point (SDP) Tunnel

Click on the Create icon **+**. Complete the fields as described below to create an SDP Tunnel. You can configure a Unicast or Multicast tunnel. Note that you can only configure a tunnel on a device configured with Loopback0 interface. If necessary, click on the **Create Loopback0** button to open the VLANs application and create a Loopback0 interface on a device. You can create either a Unicast or Multicast Tunnel. Select the corresponding button and complete the fields as described below. Click on the **Browse** button to create the tunnel on available devices. Note that only supported devices with a Loopback0 interface and the selected VXLAN Network Profile or no profile assigned will be available for selection during SDP tunnel configuration.

- **SDP Name** - The name of the SDP Tunnel.
- **Device IP (Unicast)** - The Unicast IP address of the far-end node to which customer traffic will be directed.
- **Group Address (Multicast)** - The PIM Multicast Group for the SDP Tunnel. Note that all neighbor nodes have to participate in the same multicast group to receive the VXLAN tunnel traffic from other members of the group.
- **Direction (Unicast)** - Traffic direction (Bidirectional/Unidirectional). When the service is assigned to switches, note that for Unidirectional mode, OmniVista will create SDP(s) on selected switches on the Device IP. For Bidirectional mode, OmniVista will create SDP(s) on switches selected on both the Device IP and Far End IPs. There will be one SDP created on all switches in the Far End IPs list.

- **Far End IPs** (Unicast) - The Unicast IP address(es) of the far-end node(s) to which customer traffic will be directed.
- **Device IPs** (Multicast) - The devices to include in the Multicast Tunnel. If the one or more devices selected for Multicast tunnel do not have PIM IPv4 sparse admin state enabled and PIM Bi-Direction enabled, a warning will appear listing any devices that need to be configured for PIM. Click on **Yes**, to apply the default PIM configuration profile to those devices. If you select **No**, a second warning will appear prompting you to either configure PIM on the devices or remove the devices from the configuration before proceeding.


Note: Creation of Multicast VXLAN Services requires PIM configuration on devices in the VXLAN.

Re-Applying a VXLAN Service

The **Reapply** button is enabled and can be used (in certain cases) when a VXLAN Service create/edit fails. The process of creating a VXLAN Service involves several steps (creating the SDP Tunnel, applying the VXLAN Network Profile, creating a VXLAN, binding the SDP with the VXLAN). The **Reapply** button is activated and can be used to re-apply the configuration if one of the intermediate steps fails during the process. Complete failure scenarios are not considered for re-apply. For example:

- In a Multicast Configuration - If a user applies the configuration to 3 devices, and SDP creation succeeds but VXLAN Service creation fails, the Reapply button will be enabled. But if the configuration succeeds on 2 devices completely, but the initial configuration (e.g., SDP) fails on the 3rd device, the third device will be removed from list and the Reapply button will not be enabled.
- In a Unicast Configuration - If a user configures a device and list of far end IPs, and SDP and VXLAN configuration succeed on the far end IPs but fail on the device IP, the Reapply button will be enabled. If configuration fails on the far end IP, the far end IP will be removed from list and the Reapply button will not be enabled.

Editing a VXLAN Service

Select the service and click on the Edit icon  to bring up the Edit VXLAN Service Screen. Edit the fields as described above then click on the **Save** button to save the changes to the server.

- If the edited VXLAN Service has **not** yet been assigned to switches/ports, the update will be applied and the status displayed. Click **OK** to return to the VXLAN Service Screen.
- If the edited profile **has** already been assigned to switches/ports, a confirmation prompt will appear (you can click on **Devices** to view the switches/ports). Click on the **Update** button. The update will be applied and the status displayed. Click **OK** to return to the VXLAN Service Screen.

Note: You cannot edit the Service Name. To edit the name, delete the service and re-configure the service with a new name.

Deleting a VXLAN Service

Select the service and click on the Delete icon , then click **OK** at the confirmation prompt.

SAP Profile

The VXLAN Service Access Point (SAP) Profile Screen is used to create, edit, delete, and assign SAP Profiles to switches/ports on the network. The SAP Profile is associated with a VXLAN service. A SAP ID is comprised of a customer-facing port (referred to as an access port) and an encapsulation value that is used to identify the type of network traffic to map to the associated service. You can configure up to eight (8) SAPs

per port on a switch.

Creating a SAP Profile

Click on the Create icon **+** . Complete the fields as described below and click on the **Create** button.

- **Profile Name** - A unique name for the profile (up to 32 characters)
- **VXLAN Service** - The VXLAN Service associated with the profile.
- **Trusted** - Sets the Trust Mode for the SAP Profile. If set to "True", the SAP port uses the priority value obtained from tagged packets received on the port. Untagged packets use the default port priority value. If set to "False", the priority value is set to the value configured in the Priority fields for tagged and untagged packets received on the port. (Default = Trusted)
- **Priority** - The priority value to set for tagged and untagged packets received on an untrusted SAP. (Range = 0 (lowest priority) to 7 (highest priority)).
- **Description** - A user-defined description for the SAP.

Assigning a SAP Profile


After creating a profile, select the profile and click on the Apply to Devices button to assign the profile to switches/ports on the network. Specify only ports or link aggregates that are configured as service access ports (see service access). This command does not apply to network ports.

After selecting the ports, configure the encapsulation value for the port(s) and select an Access Port Profile to associate with the SAP Profile (if applicable):


- **Encapsulation Values** - Configure the encapsulation value for the port(s). Only traffic matching this encapsulation value will be mapped to the SAP Profile.
 - **:0** - Specifies a null encapsulation value. Only untagged traffic is mapped to the profile. (Default)
 - **:all** - Specifies a wildcard SAP. All tagged traffic that is not classified into another profile is mapped to the wildcard profile.
 - **:qtag[-qtag2]** - Specifies a VLAN ID tag for ingress traffic on the access port. Only traffic with this tag is mapped to this profile.
 - **:outer_qtag.inner_qtag** - Specifies an outer VLAN ID tag and an inner VLAN tag for ingress traffic on the access port. Only double-tagged (QinQ) traffic with the specified outer and inner tags is mapped to this profile.
- **Access Port Profile** - Click on the **Browse** button and select an Access Port Profile to associate with the SAP Profile. Note that ports that have been already assigned to one Access Port Profile, cannot be assigned another profile unless all SAP profile assignments on the port are removed.

When you have completed the assignment configuration, click on the **Apply** button. The configuration will be applied and the assignment status displayed. Click **OK** to return to the SAP Profile Screen.

Editing a SAP Profile

Select the profile and click on the Edit icon  to bring up the Edit SAP Profile Screen. Edit the fields as described above then click on the **Update** button to save the changes to the server. Note that you can only edit the Priority and Description Fields. Note that if any devices are assigned to a profile, the following prompt will appear - "Update also synchronizes the changes to the device". When you click **OK**, the profile is edited and changes are synced to devices.

Deleting a SAP Profile

Select a profile and click on the Delete icon , then click **OK** at the confirmation prompt. Note that this will also delete all the corresponding SAPs on devices assigned to the profile.

Access Port Profile


The VXLAN Access Port Profile Screen is used to create, edit, and delete Access Port Profiles. An Access Port Profile is a Layer 2 Profile that is applied to an access (customer facing) port. This profile is used to specify how to process Layer 2 control frames ingressing on the access port. If an Access Port Profile is not associated with an access port, the default access profile is used to process control packets that ingress on the port. An Access Port Profile is associated with a SAP Profile using the SAP Profile Screen.

Creating an Access Port Profile

Click on the Create icon  . Complete the fields as described below and click on the **Create** button.

- **Profile Name** - A unique name for the profile (up to 32 characters)
- **VLAN Translation** - Enables/Disables egress VLAN translation for all Service Access Points (SAPs) associated with the profile. Enabling translation at the service level is only applicable if the corresponding access ports for the SAPs also have VLAN translation enabled.
- **L2 Profile Name** - A Layer 2 profile name for the Access Port Profile. You cannot use the default profile name ("default", "def-access-profile").
- **L2 Profile Attributes** - The Layer 2 attributes (e.g., STP BDU, L2 802.1x) are listed with the available behavior for each traffic type (Tunnel, Drop, Peer). The default configuration for each traffic type is pre-selected. Select the Layer 2 attributes for the profile by clicking on the desired behavior.

Editing an Access Port Profile

Select the profile and click on the Edit icon  to bring up the Edit Access Profile Screen. Edit the fields as described above then click on the **Update** button to save the changes to the server. Note that you cannot edit the Profile Name or L2 Profile Name fields. Note that if any devices are assigned to a profile, the following prompt will appear - "Update also synchronizes the changes to the device". When you click **OK**, the profile is edited and changes are synced to devices.

Deleting an Access Port Profile

Select a profile and click on the Delete icon , then click **OK** at the confirmation prompt.

VXLAN Device View

The VXLAN Device View Screen is used to view the VXLAN configuration for devices in the network. Select and option from the drop-down menu (User Switch Picker/User Topology) and click on the **Select a Device** button to select the device you want to view. The VXLAN configuration for the switch is displayed.

VXLAN Information

The following VXLAN information is displayed for the selected switch: Network Profile, VXLAN Services, SDP Tunnels, SDP Binding, SAP).

Network Profile

Displays VXLAN Network Profile information for the switch.

- **VRF Name** - The VRF instance associated with the profile.
- **UDP Port** - The UDP Port used by the VXLAN Service.
- **Loopback 0** - The Loopback 0 address for the VRF.

VXLAN Services

Displays general information about the VXLAN Services configured on the switch.

- **VNID** - The Virtual Network Identifier VNID is a 24-bit segment ID (also referred to as a VXLAN Segment ID) that is used to identify encapsulated frames. A VNID is bound to a VXLAN Service when the service is created. OmniVista will auto-generate a unique VNID for VXLAN Service if VNID is set to zero (0).
- **Service ID** - The VXLAN Service ID number.
- **Type** - The type of VXLAN Service (SPB or VXLAN is supported).
- **Description** - An optional, user-configured description for the VXLAN Service.
- **Multicast Mode** - The multicast replication mode for the VXLAN Service (Headend, Tandem, or Hybrid).
- **Admin Status** - The administrative status (Enabled/Disabled) of the VXLAN Service.
- **VLAN Translation**- The administrative status (Enabled/Disabled) of VLAN translation for all Service Access Points (SAPs) associated with the VXLAN Service. VLAN translation at the service level is only applicable if the corresponding access ports for the SAPs also have VLAN translation enabled.

SDP Tunnels

Displays the Service Distribution Point (SDP) configuration for the VXLAN Service.

- **ID** - The SDP identification number.
- **Far End ID** - The IP address (loopback0) or multicast group IP address associated with the far-end VXLAN node of the SDP.
- **Description** - An optional user-configured description for the SDP Tunnel.
- **Admin Status** - The administrative state of the SDP (Up or Down).
- **TTL** - The Time-to-Live (TTL) value for the SDP.

Binding

Displays the SDP binding configuration for the switch.

- **Service ID** - The ID number of the VXLAN Service that is bound to the SDP.
- **SDP Bind ID** - The unique SDP identification number that is bound to the VXLAN Service ID.

SAP

Displays the configuration information for the specified Service Access Point (SAP) associated with the VXLAN Service.

- **SAP ID** - The access port and encapsulation associated with the VXLAN Service.
- **Service ID** - The VXLAN Service ID number.
- **Description** - An optional description configured for the SAP. By default, the description is blank.
- **Trusted** - Whether or not the SAP is trusted (Yes or No).
- **Priority** - The 802.1p priority assigned to traffic mapped to this SAP. Applied only when the SAP is not

trusted and a priority is specified.

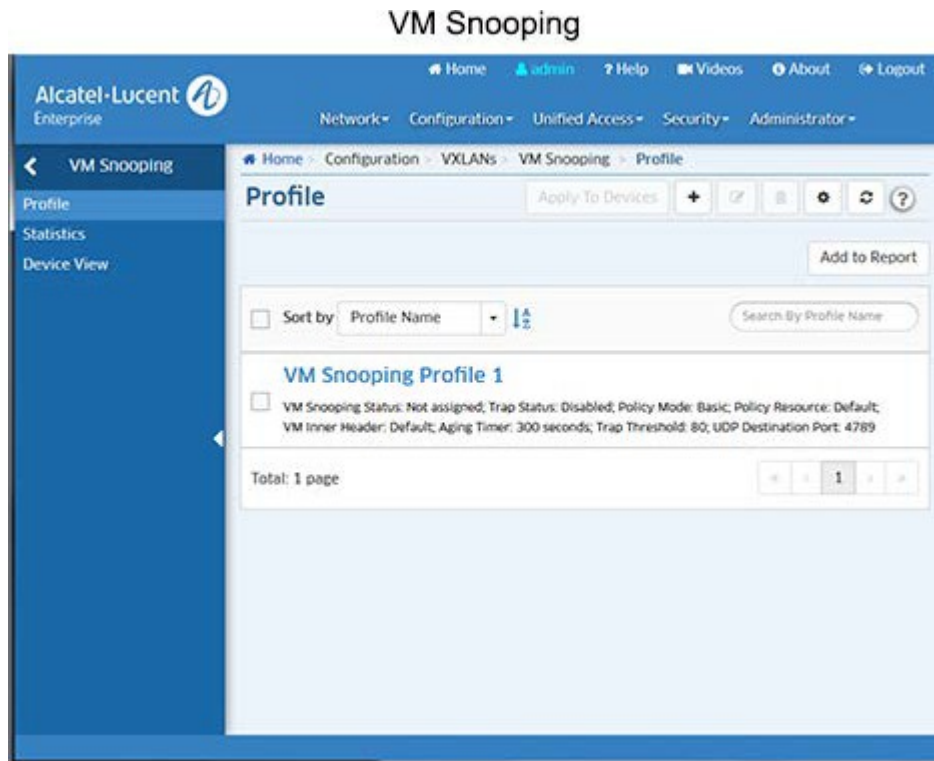
Service Access Ports

Displays Service Access Port information for the switch.

- **Port** - The slot/port number of the service access port.
- **VLAN Translation** - The administrative status of VLAN translation for the port (Enabled/Disabled).
- **L2 Profile Name** - The Layer 2 profile name for the Access Port Profile.

VM Snooping Overview

The Virtual Machine (VM) Snooping feature detects and identifies Virtual Extensible LAN (VXLAN) traffic by inspecting packets to determine if they are VXLAN encapsulated packets. Once VXLAN traffic is identified, VM Snooping collects and stores information about the VM flows in a database on the local switch. In addition to monitoring VM traffic, you can apply QoS policy list rules to the identified flows and generate SNMP traps when a new VM is learned.




To enable VM Snooping, you must create a VM Snooping Profile and assign it to switches/ports on the network. VM information can then be displayed on the VM Snooping Statistics Screen. You can also view VM Snooping Profile information for specific switches using the Device View Screen.

VM Snooping Profile

The VM Snooping Profile Screen displays all configured VM Snooping Profiles, and is used to create, edit, assign, and delete VM Snooping Profiles. A VM Snooping Profile contains global VM Snooping parameters. When a profile is created and assigned to switches/ports, VM Snooping is enabled on those switches/ports with the configured global parameters.

Creating a VM Snooping Profile


Click on the Create icon  . Enter a **Profile Name** and configure the profile as described below, then click on the **Create** button. When you are finished, select the checkbox next to the profile and click on the **Apply to Devices** button to assign the profile to switches/ports on the network.

- **Profile Name** - User-configured name for the profile.
- **Trap Status** - Enables/Disables traps for VM discovery or timeout. If Enabled, a trap is sent to the trap manager when a new VM is learned, or when a VM ages out and is removed from the system (Default = Disabled).
- **Policy Mode** - The policy lookup mode:
 - **Basic** - VXLAN UDP Port, VNI, inner source MAC, and inner IPv4 address are used for lookup (Default).
 - **Advanced** - VXLAN UDP Port, VNI, inner IPv4 source address, IP protocol, and L4 source and destination ports are used for policy lookup. In advanced IPv6 mode, VXLAN UDP Port, VNI, inner IPv6 source address, and L4 source and destination ports are used for lookup.
- **Policy Resource** - Used for configuring the hardware resources for VM Snooping:
 - **Default** - Specifies the default number of VM Policies.
 - **Extended** - Doubles the number of VM Policies.
- **VM Inner Header** - Optional inner header parameter used to specify if the header of the inner VM packet (Tagged, Untagged, Default). If you select "Default", the inner header option is set to the following values based on the current Policy Mode: Basic Mode - untagged and tagged VM packet header; Advanced Mode - tagged VM packet header.
- **Aging Timer** - The aging time value, in seconds, for VMs learned on the switch. Once a VM is discovered, it is added to the VM Snooping Database. If no flows are detected from a VM during the aging timer period, the VM is deleted from the VM Snooping Database. Information from this VM will no longer be available for display on the VM Snooping Statistics Screen (until flows are again detected and the VM is added to the VM Snooping Database. (Range = 60 - 86400, **or** 0 - *if set to "0" VMs will never age out*, Default = 300).
- **Trap Threshold** - The threshold percentage at which the switch generates a trap to indicate that VM Snooping has utilized the specified level of system resources (Range = 60 - 90, Default = 80).
- **UDP Destination Port** - The UDP destination port number(s) to look for when the switch inspects packets received on VM Snooping ports. This value is used to identify VXLAN encapsulated packets. The default port number is 4789. You can configure up to seven (7) additional UDP ports, however, configuring multiple UDP ports may slow down the VM Snooping process. Avoid using the well-known UDP ports that are already reserved by IANA for other applications.

Assigning a VM Snooping Profile

Select a profile and click on the **Apply To Devices** button. Select an option from the drop-down menu (Use Switch Picker/Use Topology and click on the **Add Remove Switch** button. Select the switch(es) to which you want to assign the profile. If VM Snooping has not already been enabled on the switch, a message will appear ("VM snooping is not enabled on ports") along with an "Add Port" link. Click on the link to bring up the port picker and select the ports to which you want to assign the profile and click **OK**. Repeat to assign the profile to additional switches/ports (select the switch, click on the **Add/Remove Ports** button, select the ports to which you want to assign the profile and click **OK**). When you have selected all of the switches/ports, click the **Apply** button. The Apply To Devices Results Screen will appear, displaying the status of the operation. Click **OK** to return to the Profile Screen.

Editing a VM Snooping Profile

Select the profile and click on the Edit icon  to bring up the Edit VM Snooping Profile Screen. Edit the fields as described above then click on the **Apply** button to save the changes to the server.

- If the edited profile has already been assigned to switches/ports, the "Update VM Snooping Profile" confirmation prompt will appear (you can click on the **Device** link to view the devices). Click **OK** to apply the update. The update will be applied and the status displayed. Click **OK** to return to the Profile Screen.

Note: You cannot edit the Profile Name. To edit the name, delete the profile and configure a new one.

Deleting a VM Snooping Profile

Select the profile and click on the Delete icon , then click **OK** at the confirmation prompt.

- If the edited profile has already been assigned to switches/ports, the "Update VM Snooping Profile" confirmation prompt will appear (you can click on **Device** link to view the devices). Click **OK** to delete the profile. The update will be applied and the status displayed. Click **OK** to return to the Profile Screen.

Removing a VM Snooping Profile From a Switch/Port

To remove a VM Snooping Profile from a switch, select the profile in the table and click on the Apply To Devices button. The switches to which the profile has been assigned will appear in the Assigned Switches area. Click on the Add/Remove Switches button to bring up the switch picker. The switches to which the profile has been assigned will appear on the right. Remove the switch(es) from the right-hand column and click **OK**. You will be returned to the Profile Screen. Click the Apply button. The configuration will be applied and the assignment status displayed. Click **OK** to return to the Profile Screen.

To remove a VM Snooping Profile from a port, select the profile in the table and click on the Apply To Devices button. The switches to which the profile has been assigned will appear in the Assigned Switches area. Select a switch and click on Add/Remove Ports button to bring up the port picker. The ports to which the profile has been assigned will appear on the right. Remove the port(s) from the right-hand column and click **OK** then click on the Apply button. The configuration will be applied and the assignment status displayed. Click **OK** to return to the Profile Screen.

VM Snooping Statistics

The VM Snooping Statistics Screen is used to display VM Snooping information for Virtual Machine (VM) traffic flows on a VM Snooping port or link aggregate. Once VM Snooping is enabled on ports, the packets flowing from the configured ports are snooped and upon matching configured UDP port, flow details are written to the VM Snooping Database. The VM Snooping Statistics Screen enables you to search for and display VM snooping information based on VM IP Address, VM MAC Address, VXLAN VNI, Destination Port, or Policy Name.

Note: OmniVista collects VM Snooping statistics from a device via an FTP session. The Telnet/FTP User Name and Password must be configured on a device for OmniVista to collect statistics. If necessary, go to the Topology application, right-click select a device(s), right click and select Edit to configure the Telnet/FTP User Name and password. If the "Prefer SSH" option is enabled in Device properties, statistics will be collected via SFTP.

Also note that by default there is a scheduler job performed every 15 minutes for collecting VM Snooping statistic data from the all supported switches (VSnoop Purge Scheduler). You can modify the interval time on the Scheduler Jobs Screen.

Searching for VM Information

You can search for VM information from a number of different sources and for different time periods. Enter the search criteria as described below, then click on the Search button. The information will be displayed in the Statistics Data Table.

- **From Date** - The start date for the VM information you want to view.
- **To Date** - The end date for the VM information you want to view.
- **Search By** - The search option to use. Select a search option, then enter the specific search criteria (e.g., VM IP Address, MAC Address).
 - **VM IP Address** - The VM IP address you want to search for.
 - **MAC Address** - The MAC address you want to search for.
 - **VXLAN VNI** - The VXLAN VNI you want to search for.
 - **Destination Port** - The destination port you want to search for.
 - **Policy Name** - The QoS Policy you want to search for.
- **Limit** - The number of rows of data to display (Range = 500 - 5,000, Default = 1,000).

VM Information

By default, all of the columns defined below appear in the Statistics Data Table. However, you can configure Custom Templates to view specific information. To configure a template, click on the Configuration icon next to one of the templates (**Custom Template 1**, **Custom Template 2**), select the column headings you want to display for that template, and click **OK**. You can configure two (2) Custom Templates. The headings you select when you configure a template will be displayed until you change them again.

- **Chassis/Slot/Port** - The physical port on which snooping is performed.
- **VTEP Source IP** - The VXLAN Tunnel end point source IP address.
- **VTEP Destination IP** - The VXLAN Tunnel end point destination IP address.
- **VXLAN VNI** - The VXLAN network identifier.
- **VTEP VLAN** - The VXLAN Tunnel end point VLAN.
- **VM Source MAC** - The source MAC of VM that is participating in the flow.
- **VM Source IP** - The source IP Address of VM that is participating in the flow.
- **VM Source Port** - The source port of the VM that is participating in the flow.
- **VM Destination MAC** - The VM destination MAC address.
- **VM Destination IP** - The VM destination IP address.
- **VM Destination Port** - The VM destination port.
- **VM IP Protocol** - The protocol that is being used by VMs in the flow (IPv4/IPv6).
- **Flow Learned Time** - The time at which the VM was identified during snooping.
- **Flow Update Time** - The most recent time verified, whether the flow is live or not.
- **Policy Name** - The name of the QoS policy rule applied to the VM flow.
- **Policy List** - The flow that is passing through by matching indicated policy list.
- **VM VLAN** - The VLAN on which the VM is learned and forwarded.
- **Sampled Packets** - The number of packets considered for snooping.

Device View

The VM Snooping Device View Screen is used to view VM Snooping Profile and port information for switches in the network. Click on the **Browse** button to select a switch then click **OK**. The VM Snooping configuration for the switch is displayed.

VM Snooping Profile Information

Displays VM Snooping Profile information for the selected switch.

- **Trap Status** - The VM discovery/timeout trap status. If Enabled, a trap is sent to the trap manager when a new VM is learned, or when a VM ages out and is removed from the system (Default = Disabled).
- **Policy Mode** - The policy lookup mode:
 - **Basic** - VXLAN UDP Port, VNI, inner source MAC, and inner IPv4 address are used for lookup (Default).
 - **Advanced** - VXLAN UDP Port, VNI, inner IPv4 source address, IP protocol, and L4 source and destination ports are used for policy lookup. In advanced IPv6 mode, VXLAN UDP Port, VNI, inner IPv6 source address, and L4 source and destination ports are used for lookup.
- **Policy Resource** - Used for configuring the hardware resources for VM Snooping:
 - **Default** - Specifies the default number of VM Policies.
 - **Extended** - Doubles the number of VM Policies.
- **VM Inner Header** - Optional inner header parameter used to specify if the header of the inner VM packet (Tagged, Untagged, Both) (Default = Both). By default, the inner header option is set to the following values based on the current policy lookup mode: Basic Mode—untagged and tagged VM packet header; Advanced Mode—tagged VM packet header.
- **Aging Timer** - The aging time value, in seconds, for VMs learned on the switch. Once a VM is discovered, it is added to the VM Snooping Database. If no flows are detected to/from a VM during the aging timer period, the VM is deleted from the VM Snooping Database. Information from this VM will no longer be available for display on the VM Snooping Statistics Screen (until flows are again detected and the VM is added to the VM Snooping Database. (Range = 0 - 86400, Default = 300). If set to "0", VMs will never age out.
- **Trap Threshold** - The threshold percentage at which the switch generates a trap to indicate that VM Snooping has utilized the specified level of system resources (Range = 60 - 80, Default = 80).
- **UDP Destination Port** - The UDP destination port number(s) to look for when the switch inspects packets received on VM Snooping ports. This value is used to identify VXLAN encapsulated packets. The default port number is 4789. You can configure up to seven (7) additional UDP ports, however, configuring multiple UDP ports may slow down the VM Snooping process. Avoid using the well-known UDP ports that are already reserved by IANA for other applications.
- **Hardware Allocation Status** - The hardware resource allocation status for VM Snooping.

Enabled VM Snooping Port Information

Displays basic information for VM Snooping-enabled ports for the selected switch.