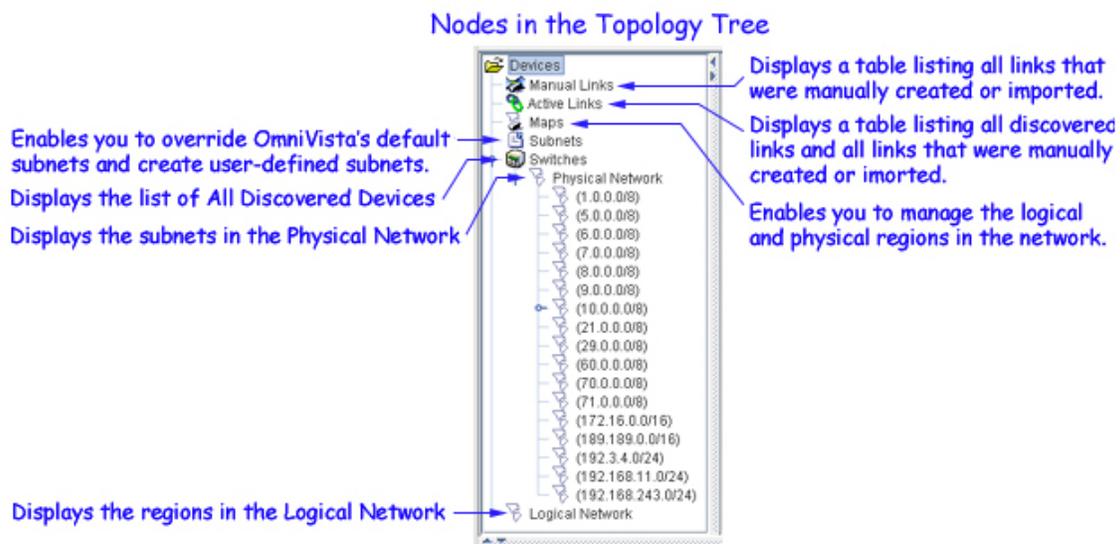


Getting Started with Topology

The Topology application enables you to manage the physical devices in the network and view the topology of the network. All physical devices in the network are listed in the Topology application's list of All Discovered Devices and in the Topology Tree. Popup menus in the list of All Discovered Devices and the Tree provide the functionality needed to manage and configure devices. Clicking on an individual device in the Topology Tree connects you to the device.

The Topology application enables you to view the devices and links in the network in various ways. You can display the overall physical network to view its subnets and the links between them. You can display individual subnets and the individual devices therein. You can create maps of "logical" regions that enable you to group and display devices in a way that is meaningful for your individual network configuration.

Nodes in the Topology Tree enable you to view tables listing all network links, manage the logical and physical regions in the network, create user-defined subnets in the Physical Network, view the list of All Discovered Devices, and view graphical maps of the subnets in the Physical Network and the regions in the Logical Network. Each node in the Topology Tree is described below.



Manual Subnets

OmniVista now provides the ability to create manual (i.e., user-defined) subnets. In previous releases of OmniVista, subnets were automatically created by default. The new Subnets node in the Topology Tree enables you to override OmniVista's default subnet creation and manually define the subnets that OmniVista displays in the Tree. If manual subnets exist when a discovery is performed, OmniVista will place the discovered switches into the manual subnets upon their discovery. If manual subnets are created after discovery, OmniVista will place known switches into the manual subnets when they are created. Click [here](#) for more information on both user-defined subnets and OmniVista's default subnet creation.

Subnet Labels

In the Tree, subnets are labeled in the form *ipaddress/n*. The */n* indicates the number of bits in *ipaddress*, starting from the left, that identify the network (i.e., the subnet). These bits will have the same value in all addresses that belong to the subnet. The literal value of these bits displays in *ipaddress*. Any bits in *ipaddress* that do not identify the subnet are represented by zeros.

For example, the screen above shows a subnet named **10.255.11.0/24**. The **/24** means that the first 24 bits of the address, starting from the left, identify the subnet and will be common to all address in the subnet. The literal value of these 24 bits, 10.255.11, displays in the subnet name. The last bits are represented by a 0, as these bits do not identify the subnet. (They identify devices.) This subnet could also be represented as 10.255.11.*, where the * character represents any value. This subnet will include all devices with an IP address in the range 10.255.11.0 - 10.255.11.255.

As a second example, consider a subnet named **10.0.0.0/8**. The **/8** means that the first eight bits of the address identify the subnet and will be common to all address in the subnet. The literal value of these eight bits, 10, displays in the subnet name. All other bits are represented by zeros. This subnet could also be represented as 10.*.*, where the * character represents any value. This subnet will include all devices with an IP address in the range 10.0.0.0 - 10.255.255.255.

The List of All Discovered Devices and the Tree

All devices discovered display in the list of All Discovered Devices. Select **Switches** in the Tree to view the list of All Discovered Devices, as shown below. Click here for information on the fields in the list of All Discovered Devices. Each discovered device also displays in the Tree. Click **Switches** and **Physical Network** open in the Tree, as shown below, to display the individual subnets in the Physical Network. Click a subnet open to view the individual devices on the subnet. Color coding in the Tree and in the list of All Discovered Devices provides information on the state of each device.

Select Switches in the Tree to display the list of All Discovered Devices

Discovered devices display in the list of All Discovered Devices

Click Switches and Physical Network open in the Tree to list the subnets in the network. Click a subnet open to view individual devices.

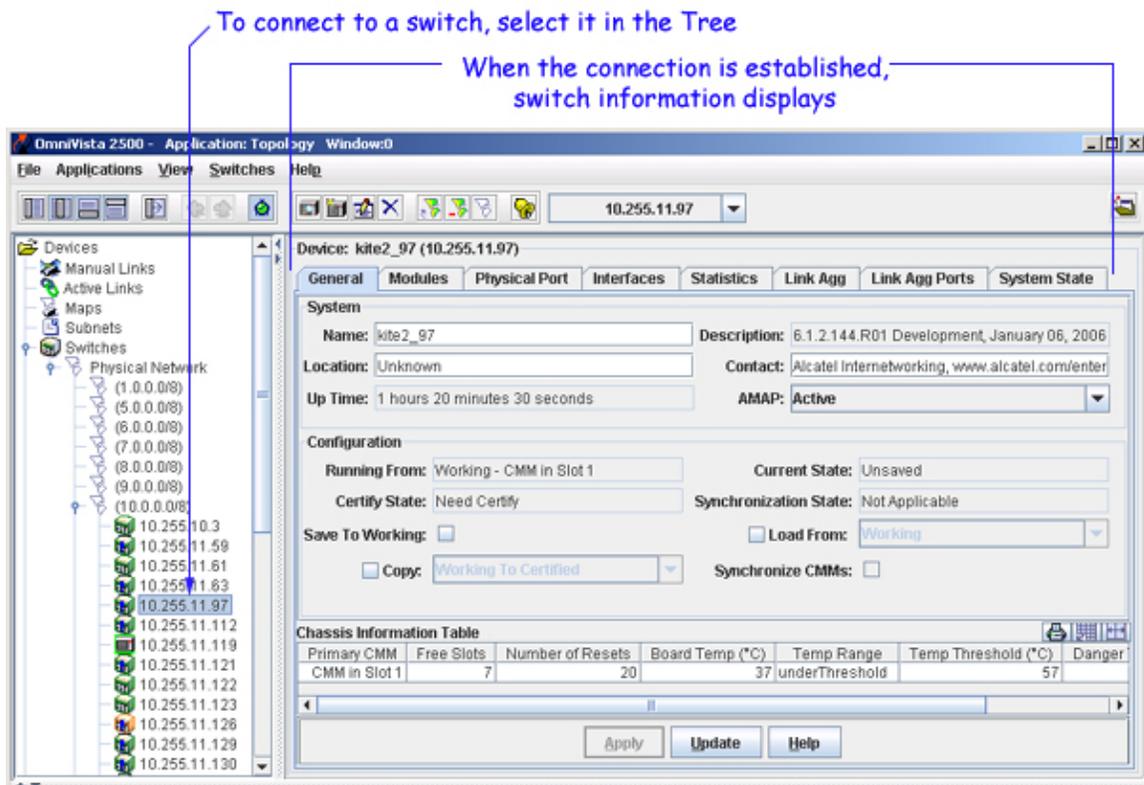
| Name | Address | DNS Name | Type | Version |
|--------------------|---------------|----------|-----------|---------------|
| DCTestnetCore | 10.255.10.3 | | OS7700 | 5.1.6.96.R02 |
| Kite_59 | 10.255.11.59 | | OS6800-48 | 6.1.2.141.R01 |
| wTarget | 10.255.11.61 | | OS6800-24 | 6.1.3.50.R01 |
| wTarget | 10.255.11.63 | | OS6800-24 | 6.1.2.140.R01 |
| kite2_97 | 10.255.11.97 | | OS6800-24 | 6.1.2.144.R01 |
| Kite | 10.255.11.112 | | OS6800-24 | 6.1.2.144.R01 |
| no-name-119 | 10.255.11.119 | | OmniSIR-5 | 4.5.2 |
| wTarget | 10.255.11.121 | | OS9700 | |
| WV_HAWK_122 | 10.255.11.122 | | OS6648 | 5.4.1.163.R01 |
| NMS_123_Hawk | 10.255.11.123 | | OS6648 | 5.1.5.133.R04 |
| WV_FUJII_126 | 10.255.11.126 | | OS9700 | 6.1.1.633.R01 |
| WV_FUJII_129 | 10.255.11.129 | | OS9800 | 6.1.1.633.R01 |
| ES0001A-1 | 10.255.11.130 | | OS6648 | 5.1.6.19.R03 |
| Kite_135 | 10.255.11.135 | | OS6800-48 | 6.1.2.140.R01 |
| NMS-test-148 | 10.255.11.148 | | Omni-SWX | 4.5.3.100 |
| AOS_Hawk_157_alias | 10.255.11.157 | | OS6624 | 5.1.6.140.R02 |
| Alias | 10.255.11.174 | | Omni-SWX | 4.4.5 |
| nms189 | 10.255.11.189 | | OS8800 | 5.1.6.164.R02 |
| OmniStack 8008_201 | 10.255.11.201 | | OS-8008 | V2.50.09 |

Adding Devices Manually

It is possible to add devices manually to the list of All Discovered Devices using the New Discovery Manager Entry window. It is also possible to import a list of devices from a .csv file to the list of All Discovered Devices. You can also export the list of All Discovered Devices to a .csv file (where it can be edited).

Connecting to a Switch

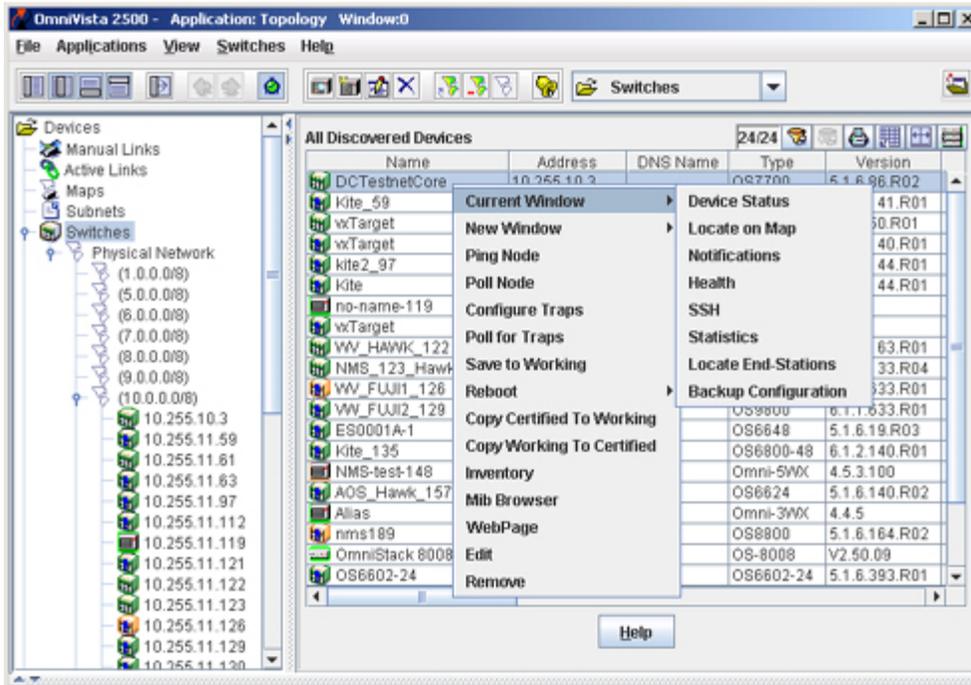
You can connect to a switch merely by selecting it in the Tree. When the connection is established, tabs of information on the switch display, as shown below. Note that the information displayed is somewhat different for AOS devices (the OmniSwitch 6000/7000/8000/9000 Product Series), various XOS devices (early generation OmniSwitch, OmniStack, OmniAccess 512, and OmniSwitch/Router devices), and third-party devices. Click the **Help** button at the bottom of any tab for specific information on the fields in each tab.



Popup Menu Functionality

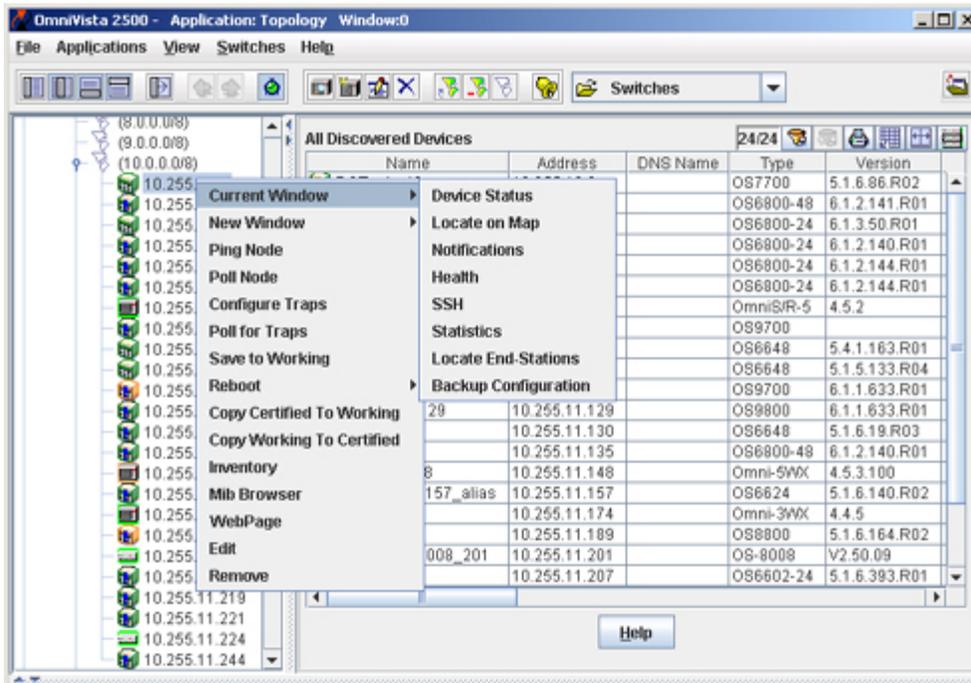
You can click right on one or more devices in the list of All Discovered Devices to display a popup menu. Somewhat different versions of the popup menu display for AOS devices, XOS devices, and third-party devices. The popup menu for AOS devices is shown below. Note that several menu items on the popup menu are active when multiple switches are selected. This enables you to perform the respective function on multiple switches simultaneously. Click here for specific information on the menu items on the popup menu.

Popup Menu for AOS Devices
(OmniSwitch 6000/7000/8000 Series)
Right-click on any device in the list to display the menu.



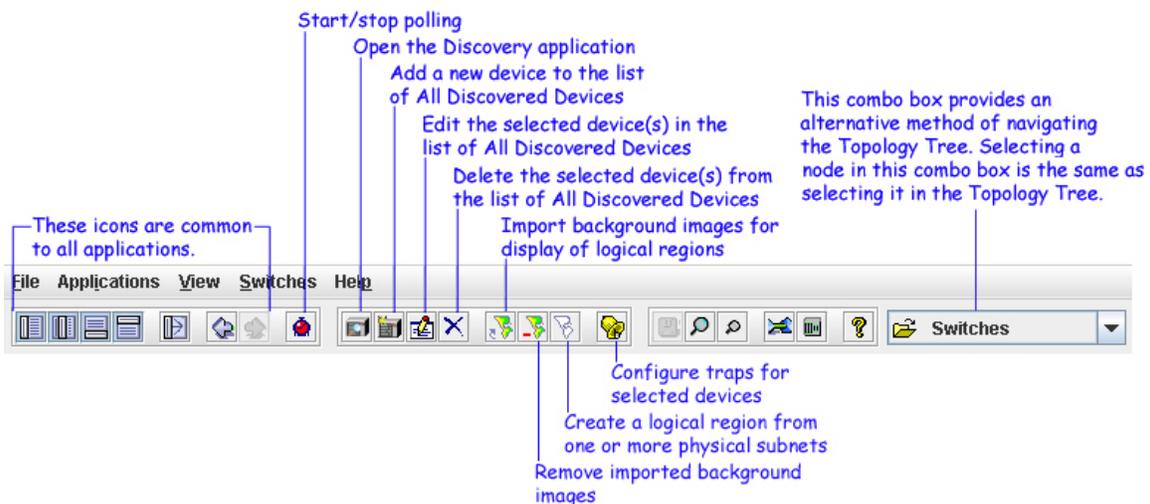
You can also click right on any device listed in the Tree to display a popup menu. Somewhat different versions of the Tree popup menu display for AOS devices, XOS devices, and third-party devices. The Tree popup menu for AOS devices is shown below. All menu items on the Tree popup menu also appear on the popup menu in the list of All Discovered Devices (described above).

Tree Popup Menu for AOS Devices
 Right-click on any device in the tree to display the menu.



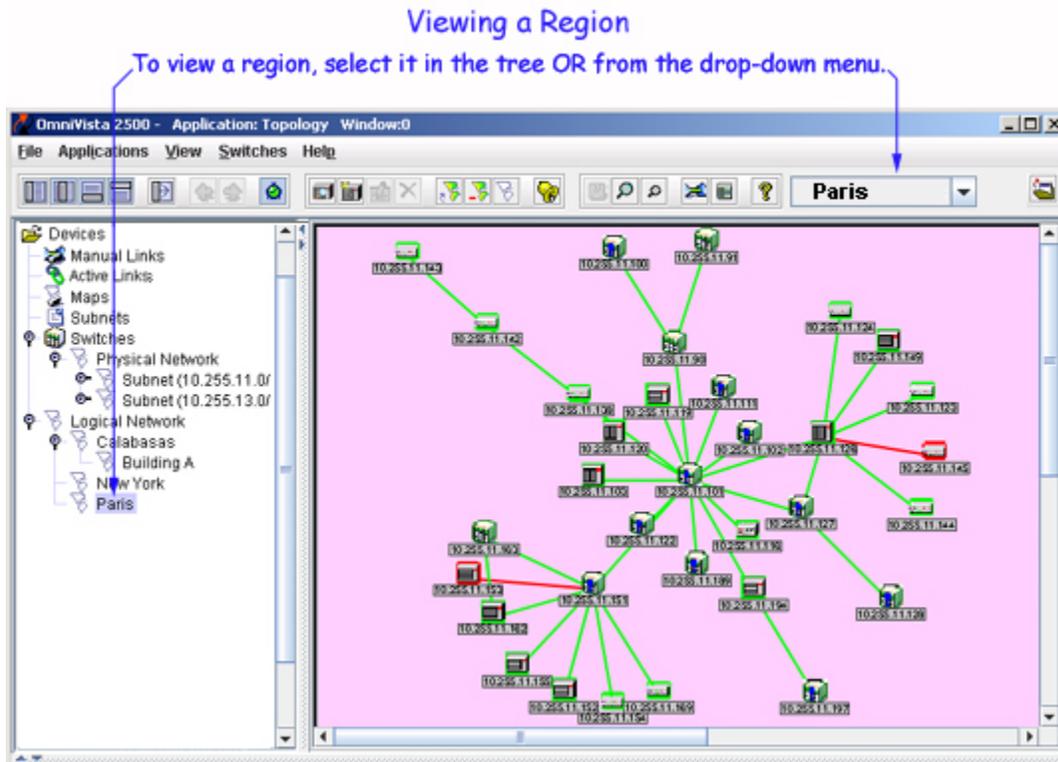
The Topology Toolbar

The toolbar that displays when the Topology application opens contains tools that enable you to perform specific tasks quickly, as shown and explained below.



Viewing the Network

The Topology application enables you to display and view the topology of any network region, including the overall Physical Network, the overall Logical Network, or any individual subnet or region therein. Color coding in the display provides status information on each region, device, and link displayed. Specific information about the links in each region can be viewed. Popup menus provide further functionality. To view any network region, select it in the Tree or in the combo box shown below. Network regions display with the background color and background image specified when the region was created or edited. Click [here](#) for detailed information on the functionality available when viewing the network.



The Physical Network and the Logical Network

The Physical Network, as its name implies, is an image of the physical subnets and devices in the network. When OmniVista discovers the network, it arranges the discovered devices into default subnets. You can override OmniVista's default subnet creation by creating manual, that is, user-defined, subnets. (Click [here](#) for more information on default and manual subnet creation.) However, all subnets in the Physical Network, both default subnets and manual subnets, are created according to the device IP address. You cannot "pick and choose" the individual devices to be included in a subnet.

In contrast, within the Logical Network you can create "logical regions" and select the individual devices to be included in the region, regardless of the device IP address. You can create logical regions where devices are grouped and displayed in any way that is meaningful for your individual network, in any configuration desired.

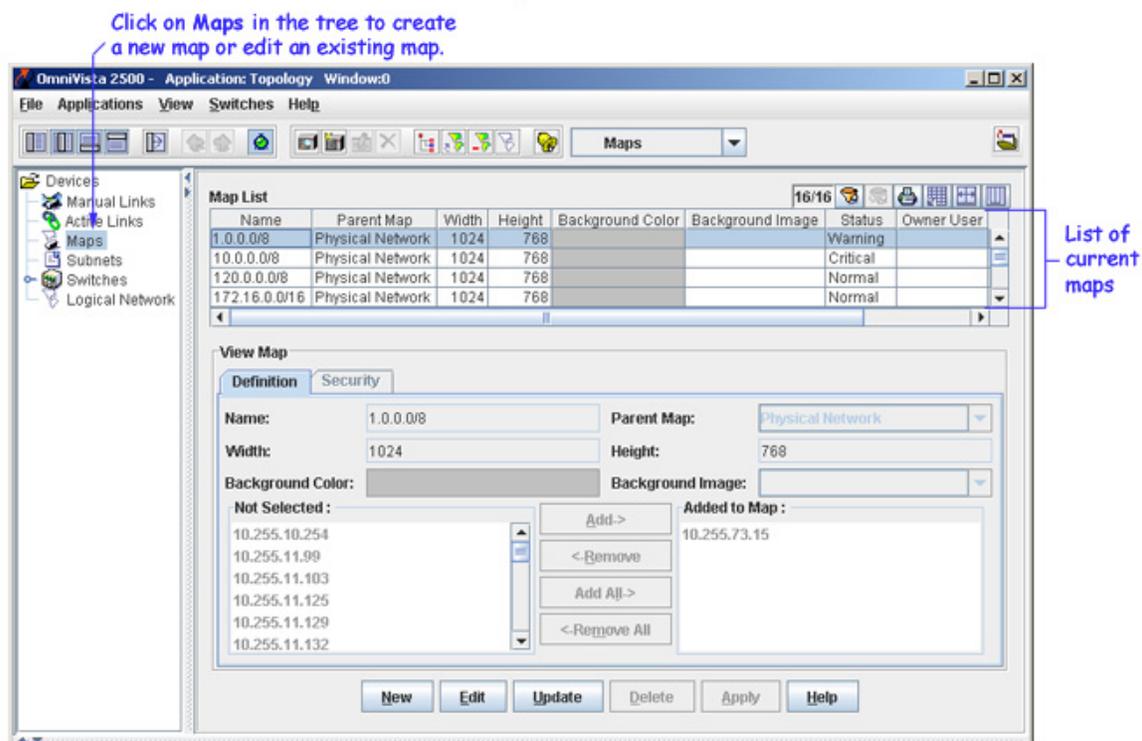
When you create a regional map in the Logical Network, you must define a "parent" map. The parent map can be the Logical Network itself, as it is for Calabastas, New York, and Paris in the screen shown above.

You can also "nest" regional maps in the Logical Network by specifying an existing regional map as the parent map. For example, Calabasas is the parent map of Building A in the screen shown above.

Network Maps

The Maps window, shown below, enables you to create regional maps in the Logical Network from scratch, create regional maps in the Logical Network from existing subnets in the Physical Network, edit existing regional maps in both the Logical Network and the Physical network, and delete regional maps from the Logical Network or the Physical Network. When you create or edit a regional map, you can define the background color you want used when the map is displayed and the width and height of the viewing window. You can also specify a background image for the map, if desired. Click here for detailed information on creating and managing regional maps.

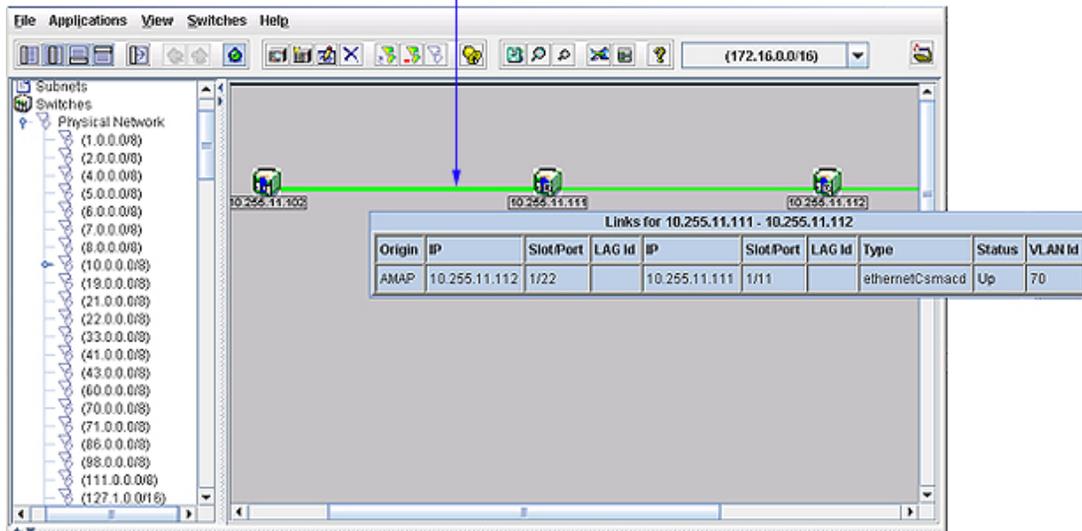
The Maps Window



Network Links

Both the ATM and Ethernet links in the network can be automatically discovered during the discovery process. When a regional map is displayed, the links in the region also display and are color-coded as to their status. Whenever you are viewing a regional map, you can display information about the links in the region. To do this, merely place the cursor on the connection you wish to view. A table listing the individual links in the connection displays, as shown below. Click here for more information on viewing network links.

Place the cursor on a connection you want to view. A table listing all links in the connection displays.



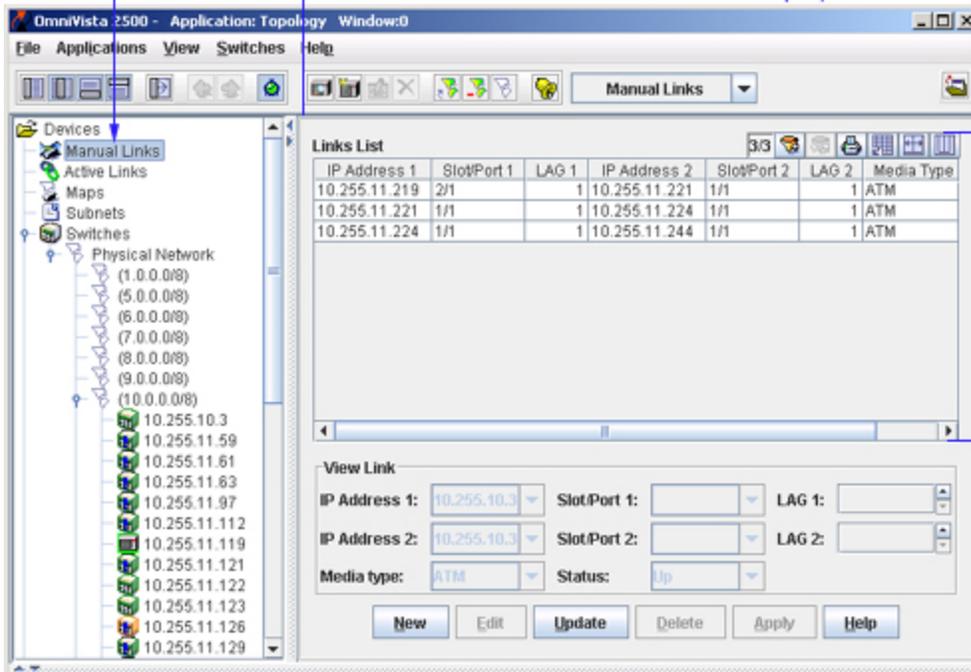
In addition to discovering links via the discovery process, you can create links "manually" in OmniVista via the Manual Links window, shown below. You can also import a list of links from a .csv file. Alternatively, you can export a list of manual links to a .csv file (where it can be edited). Click here for detailed information on creating and managing manual links.

The Manual Links window displays only links that were created manually or imported into OmniVista. All network links, including discovered links, are displayed in the list of Active Links. Note that discovered links cannot be listed, edited, or deleted.

The Manual Links Window

Click Manual Links in the Tree

The Manual Links window displays

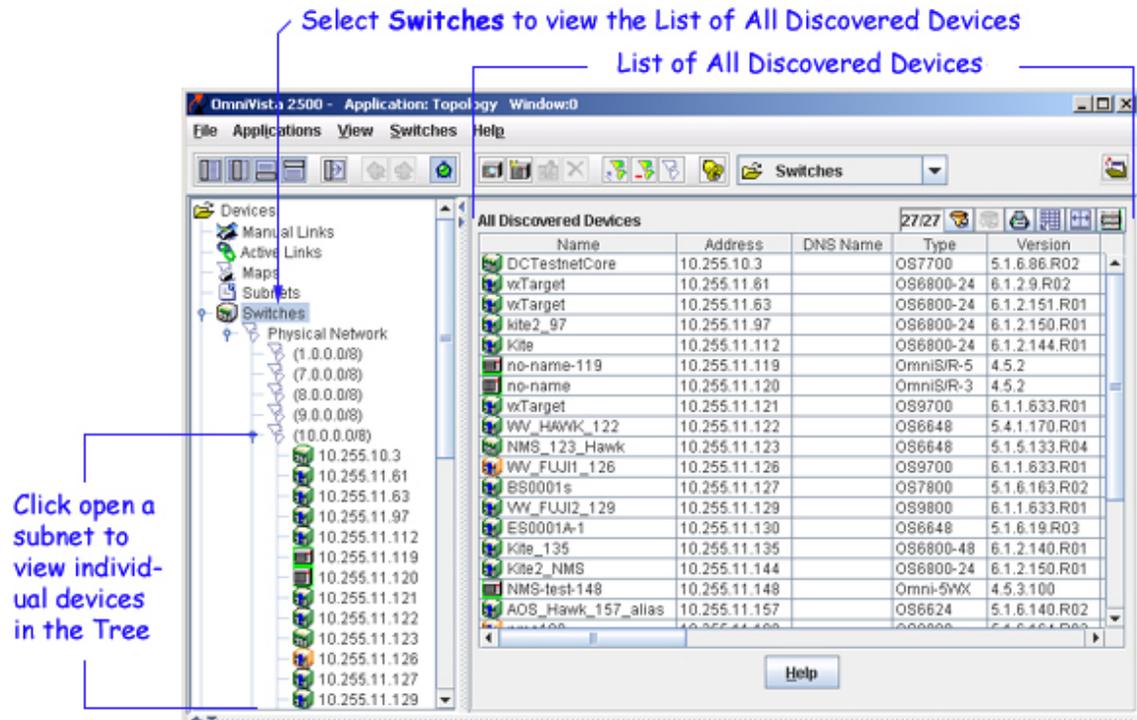


Manually created links and imported links are listed here

The List of Discovered Devices and Devices in the Tree

The list of All Discovered Devices is a list of all devices that were discovered, and all devices that were added to the list manually. To display the list of All Discovered Devices, select **Switches** in the Topology tree, as shown below. Information on each device in the list is provided in tabular form. In addition, the list of All Discovered Devices enables you to perform functions on a single switch or on multiple switches simultaneously. To do so, simply select a single device in the list, or select multiple devices in the list, and then click right to display a popup menu of the functions available.

The Physical Network in the tree lists each known subnet. You can click open a subnet in the tree to view the individual devices. You can connect to any device merely by selecting it in the tree. Popup menus available in the tree provide additional functionality. You can only select one switch at a time in the tree.



Color Coding

Entries in the list of All Discovered Devices and device icons in the tree can display green, red, or orange. Devices displayed in green are up (responding to OmniVista's polls). Devices displayed in red are down (not responding to OmniVista's polls). Devices displayed in orange are in the warning state (the switch has sent at least one warning or critical trap).

In addition, icons for AOS devices display a blue exclamation mark () when the switch configuration is in the Unsaved state (changes have been made to the running configuration of the switch that have not been saved to the working directory) or the Uncertified state (the working directory has changes that are not in the certified directory). Click here for more information.

Adding a New Device Manually

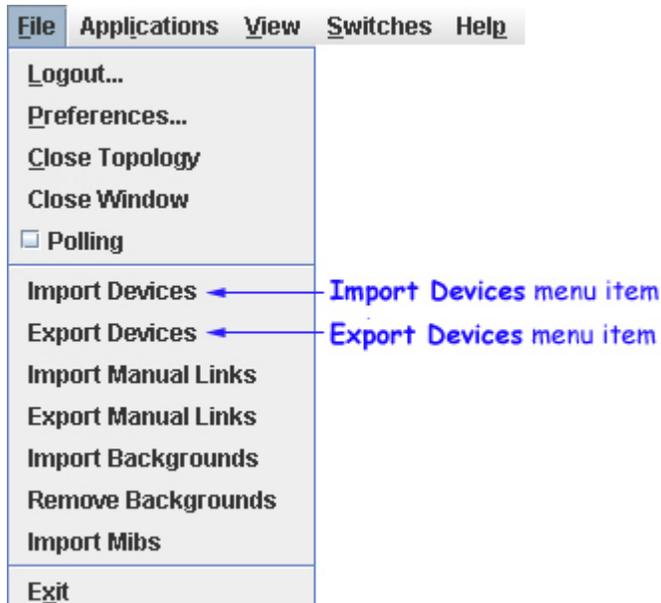
You can add a new device to the list of All Discovered Devices and the tree manually. To do this, click the "New Switch" icon, shown below, or select **New** on the **Switches** menu. In either case the New Discovery Manager Entry window displays, which enables you to add a new device. Click here for more information.



Importing and Exporting Devices

You can import devices into the list of All Discovered Devices from a Microsoft Excel file or any other application that produces comma-separated value files (.csv file extension). A comma-separated value file, as the name implies, lists a series of values separated by commas. To import a list of devices, select **Import Devices** on the File menu, shown below. The Import Devices window displays, which enables you to locate the .csv file you want to import. Locate the file and then click the **Import** button on the window. All imported devices display in the list of All Discovered Devices.

You can export devices to a .csv file (which can be edited) in much the same manner: select **Export Devices** on the File menu to display the Export Devices window, navigate to the location where you want to save the .csv file, and then click the **Export** button on the window. All devices on the list of All Discovered Devices are saved in the .csv file.



Using the List of All Discovered Devices

The following section describes the information fields in the **All Discovered Devices** table.

Information Fields in the List

Name

The name of the device.

Address

The address of the device.

DNS Name

The DNS name of the device.

Type

The type of the device chassis.

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.

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.

In all "Failed" cases, "Reload From Working" will be disabled on the switch until a successful upgrade is performed.

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.

Status

This field displays the operational status of the device. It displays **Up** if the device is up and responding to polls. (When a device is up, it displays green in both the List of All Discovered Devices and the tree.) It displays **Down** if the device is down and not responding to polls. (When a device is down, it displays red in both the List of All Discovered Devices and the tree.) This field displays **Warning** if the switch has sent at least one warning or critical trap and is thus in the warning state. (When a device is in the warning state, it displays orange in both the List of All Discovered Devices and the tree.)

Traps

This field indicates 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 switch. OmniVista will read the switch's trap configuration when traps are configured for the switch via the Configure Traps Wizard.

Seen By

This field lists the Security Groups that are allowed to view the device. (The Security Groups that are allowed to view a device can be defined when devices are autodiscovered, added manually, or edited.) The default Security Groups shipped with OmniVista are as follows:

Default group. This group has read-only access to switches in the list of All Discovered Devices that are configured to grant access to this group.

Writers group. This group has both read and write access to switches in the list of All Discovered Devices that are configured to grant access to this group. However, members of this group cannot run autodiscovery nor can they manually add, delete, or modify entries in the list of All Discovered Devices.

Network Administrators group. This group has full administrative access rights to all switches on the network. Members of this group can run autodiscovery and can manually add, delete, and modify entries in the list of All Discovered Devices. Members of this group also have full read and right access to entries in the Audit application and the Control Panel application. Members of this group can do everything EXCEPT make changes to Security Groups.

Administrators group. This group has all administrative access rights granted to the Network Administrators group AND full administrative rights to make changes to Security Groups.

Note that other Security Group names may display in this field if custom Security Groups were created. Refer to help for the Security application *Users and Groups* for further information on Security Groups.

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. (Note that you can specifically command a switch to reboot from either directory -- [click here for information.](#))

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. [Click here for more information.](#)

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:

Unsaved. Changes have been made to the running configuration of the switch that have not been saved to the working directory.

Uncertified. Changes have been saved to the working directory, but the working directory hasn't been copied to the certified directory. The working directory and the certified directory are thus different.

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.

OmniVista is now capable of tracking AOS configuration changes made through CLI commands or WebView, and so will reflect configuration changes made outside of OmniVista through these two interfaces in the Changes field. 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. Whenever an AOS device is in the Unsaved or Uncertified state, a blue exclamation mark displays on its icon ().

Discovered

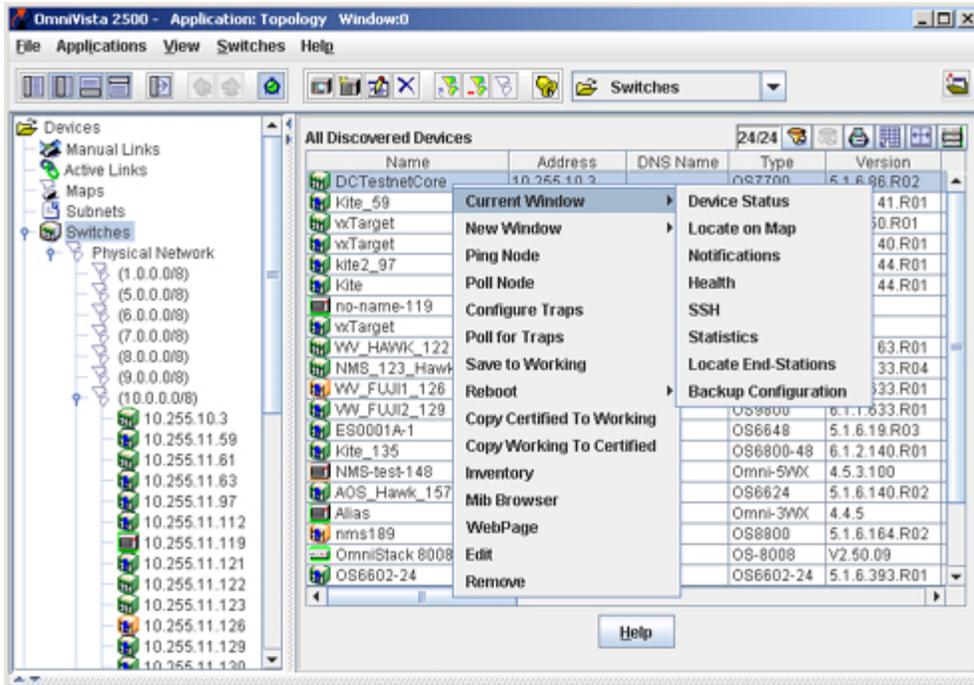
This field displays the date and time when OmniVista successfully pings or polls the switch for the first time. This value remains unchanged until the switch entry is deleted. This field will remain blank if OmniVista does not ping or poll the switch at all.

Popup Menu in the List

Click right on one or more devices in the list of All Discovered Devices to display a popup menu. Somewhat different versions of the popup menu display for various AOS devices, XOS devices, and third-party devices. The popup menu for AOS devices is shown below. Note that several menu items on the popup menu are active when multiple switches are selected. This enables you to perform the respective function on multiple switches simultaneously.

The first two items on the popup menu, **Current Window** and **New Window**, each expand to multiple menu items. **Current Window** and **New Window** enable you to open their respective menu items in the current OmniVista window or in a new, additional OmniVista window. Each individual menu item that can display on the popup menu is explained below.

**Popup Menu for AOS Devices
(OmniSwitch 6000/7000/8000 Series)**
Right-click on any device in the list to display the menu.



Current Window or New Window > Device Status

Selects the switch in the Tree and establishes a connection to the switch, exactly as if you had manually selected the switch in the Tree. If the switch's icon is not visible in the Tree, OmniVista will expand the Tree and scroll until the switch icon is visible. When a connection is established, device-specific configuration and statistics information displays. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Locate on Map

Loads and displays a regional map in the Physical Network that contains the selected device. The device is automatically selected and centered in the map display. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Notifications

Loads the Notifications application for the selected switch. The Notifications application enables you to view traps for the switch. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Health

Loads the Health application for the selected switch. The Health application displays information on the health of the selected switch. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Telnet or SSH

Either **Telnet** or **SSH** (Secure Shell) displays by default on the popup menu, as user-configured for the individual switch. You can configure the default selection for a switch through any of the methods described below. You can also define the switch's Telnet user name and password to OmniVista by means of these methods. When the Telnet user name and password are known, OmniVista will auto login for your

convenience when Telnet or SSH sessions are established. Configure the defaults for a switch using any one of the following methods:

Discover the switch with an SNMP setup that has its **Shell Preference** field set to **Telnet** or **SSH**, as desired. Enter the Telnet user name and password in the respective fields on the SNMP Setups window. (For more information, refer to the help for the Discovery application.)

Edit the switch after discovery and activate the **Prefer SSH** checkbox on the General Tab of the Edit Discovery Manager Entry window. This will specify that SSH is the default for the switch. Enter the Telnet user name and password in the respective fields.

Activate the **Prefer SSH** checkbox on the New Discovery Manager Entry window when you add a switch manually. This will specify that SSH is the default for the switch. Enter the Telnet user name and password in the respective fields.

The **Telnet** or **SSH** menu item opens the Telnet application and establishes a Telnet or SSH connection, respectively, with the selected switch. If the switch's Telnet user name and password are known to OmniVista, auto login will occur. Otherwise you will need to manually enter the switch's Telnet user name and password. Each time the **Telnet** or **SSH** menu item is selected, a new Telnet or SSH session is established. Individual Telnet and SSH sessions are identified by tabs that display the switch IP address. Telnet or SSH sessions can be established in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Statistics

Loads the Statistics application with the Add Item window open and the relevant switch selected automatically. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Locate End-Stations

Loads 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 tab. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Backup Configuration

Loads the Backup Configuration utility in the Health application for the selected switch. The Backup Configuration utility in the Resource Manager application loads and saves firmware files for the selected switch. This function can be performed in the current OmniVista window or in a new OmniVista window.

Ping Node

Causes an immediate ping to the selected switches. The result of the ping -- an "equipment is alive" message or an "equipment does not respond" message -- is reported in the Status Panel.

Poll Node

Causes an immediate poll of the selected switches. The success or failure of the poll is reported in the Status Panel.

Configure Traps

Opens the Configure Traps Wizard for the selected switches. The Configure Traps Wizard enables you to configure traps for the switches.

Poll for Traps

Causes an immediate poll of the selected switches for traps. The success or failure of the poll is reported in the Status Panel. Traps are reported in the Notifications application. You can also manually poll for traps from a *single* switch by right-clicking on the switch in the Device Tree in Topology, or any OmniVista application displaying a Device Tree (e.g., VLANs, Notifications).

Save to Working (AOS Devices)

Saves the primary CMM's current running configuration to the working directory of the switch. Executing this command is the same as executing the Save To Working command for an individual device. However, when the List of All Discovered Devices is displayed, the **Save to Working** menu item enables you to save the configurations of multiple switches in one operation. Click here for an overview of the directory structure of AOS devices.

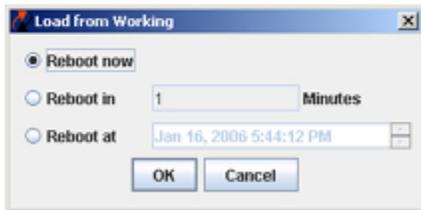
Note: When you apply the **Save to Working** option on a device(s), you must allow 120 seconds of time to elapse, before you perform the same again.

Reboot > From Working (AOS Devices)

Reboots the primary CMM from the working directory. Executing this command is the same as executing the Load From Working command for an individual device. However, when the List of All Discovered Devices is displayed, the **Reboot** menu item enables you to reboot the primary CMMs in multiple switches in one operation. Note that any unsaved configuration changes will be lost: you can save configuration changes with the **Save to Working** command before executing **Reboot**. Click here for an overview of the directory structure of AOS devices.

When you select **Reboot > From Working**, the Load from Working window displays. The Load from Working window is shown below. This window enables you to specify whether you wish to reboot immediately (**Reboot now**), or reboot within 1 - 1000 minutes (**Reboot in x Minutes**), or reboot at a specified date and time (**Reboot at date time**). Specify the desired reboot time and then click the **OK** button.

The Load from Working window enables you to schedule the reboot.



Reboot > From Certified (AOS Devices)

Reboots the primary CMM from the certified directory. Executing this command is the same as executing the Load From Certified command for an individual device. However, when the List of All Discovered Devices is displayed, the **Reboot** menu item enables you to reboot the primary CMMs in multiple switches in one operation. Note that any unsaved configuration changes will be lost: you can save configuration changes with the **Save to Working** command before executing **Reboot**. Click here for an overview of the directory structure of AOS devices.

When you select **Reboot > From Working** or **From Certified**, the Load from Certified or Load from Working window displays, respectively. The Load from Certified window is shown below. This window enables you to specify whether you wish to reload an entire switch (**Reload Entire Switch**), reboot immediately (**Reboot now**), or reboot within 1 - 1000 minutes (**Reboot in x Minutes**), or reboot at a specified date and time (**Reboot at date time**). Specify the desired reboot time and then click the **OK** button.

The Load from Certified window enables you to schedule the reboot.



Note: When you reboot the primary CMM from the certified directory, the switch will automatically failover to the secondary CMM (in other words, the two CMMs will trade primary and secondary roles). When you reboot the primary CMM from the working directory, no failover occurs.

Copy Certified to Working (AOS Devices)

Copies the contents of the certified directory in the primary CMM to the working directory in the primary CMM. Executing this command is the same as executing the Copy Certified to Working command for an individual device. However, when the List of All Discovered Devices is displayed, the **Copy Certified to Working** menu item enables you to copy the contents of the certified directory to the working directory in multiple CMMs in one operation.

Copy Working to Certified (AOS Devices)

Copies the contents of the working directory in the primary CMM to the certified directory in the primary CMM, in a manner similar to the **Copy Certified to Working** command described above.

Note: The **Copy Working to Certified** command also automatically synchronizes the switch's CMMs after the copy operation is completed.

Inventory

Loads the Inventory application for the selected switches. The Inventory application enables you to create reports. The reports can include system information, detailed module information, chassis information, and health information.

MIB Browser

Loads the OmniVista MIB Browser for the selected switch.

Note: If a read-only user launches MIB browser of a switch which is configured to use SNMPv3, the username/password specified by the OV administrator for SNMPv3 is ignored, and is substituted by "public" for the user name, authentication password, and privacy password which means that such an account must pre-exist on the switch.

WebPage or SwitchManager or TrackView

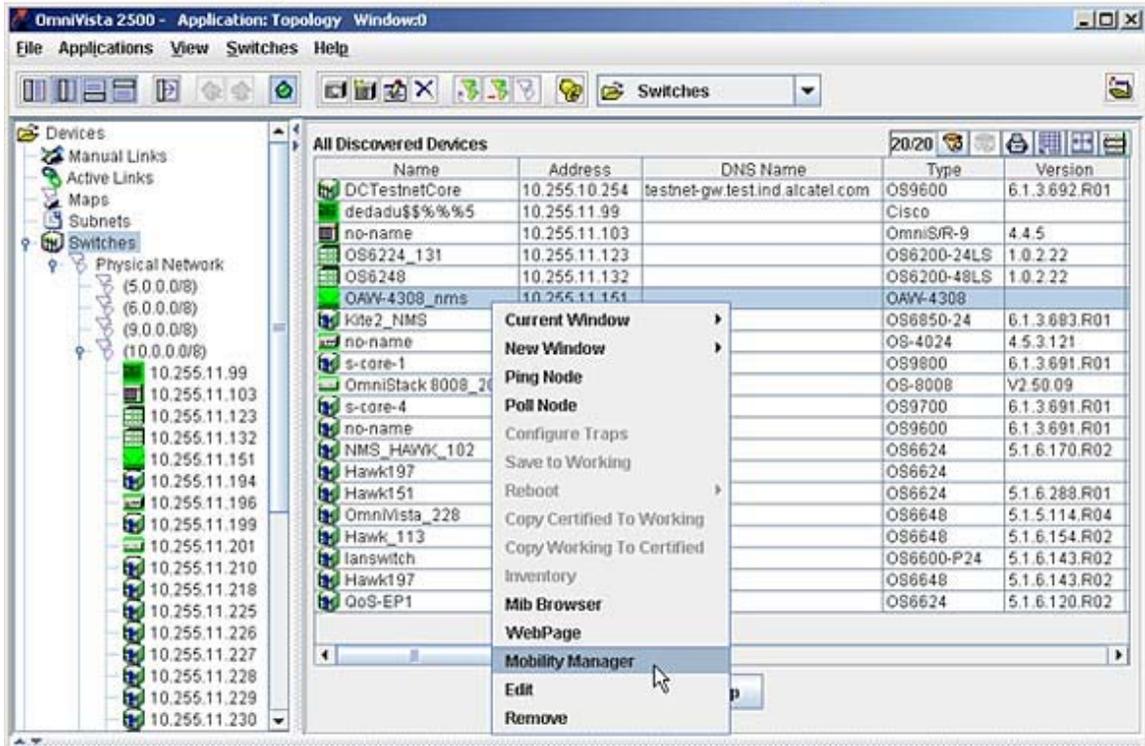
This menu item opens the device manager that is appropriate for the selected switch. WebView, the Alcatel device manager, opens for AOS devices. WebView enables you to perform direct device-level AOS configuration from a browser. **TrackView** opens for OmniCore devices. **WebPage** opens for the OmniStack 1024, 6024, 6300-24, and 8008, as well as the OmniMSS. **SwitchManager** opens for all other XOS devices. Each device manager enables you to perform device-level configuration of the selected device.

Note: SwitchManager and TrackView will open only if the respective program is installed on the client.

Mobility Manager

Launches the OmniVista default browser with a URL pointing to the Mobility Manager application for the selected wireless switch.

Popup Menu for AOS Devices
(OmniSwitch 6000/7000/8000 Series)
Right-click on any wireless device in the list to display menu.



The Mobility Manager URL can be set using the **Mobility Manager URL** option in the Preferences application. However, if the Mobility Manager URL is not defined in Preferences, then you will be prompted to define the URL in the **Mobility Manager URL** dialog box (shown below) when you select the **Mobility Manager** menu item for the selected wireless switch.



Edit

Opens the Edit Discovery Manager Entry window, which enables you to edit devices in the List of All Discovered Devices. When you edit a device, it is important to understand that you are editing OmniVista's knowledge of the device, not the device itself. Click here for the reasons you might want to edit a device. Note that you can edit multiple devices simultaneously.

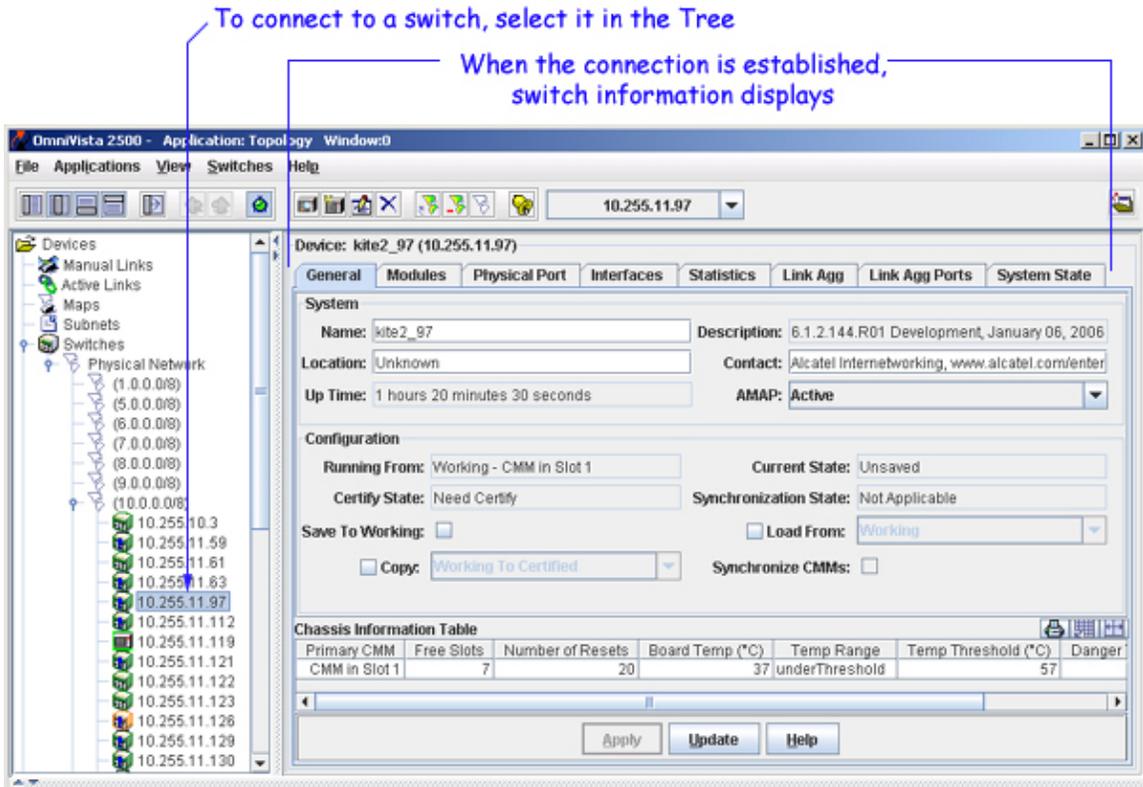
Remove

Deletes the selected devices from the list of All Discovered Devices and from the Physical Network. When a device is removed, OmniVista no longer has knowledge of the device.

Using the Tree

Connecting to a Switch

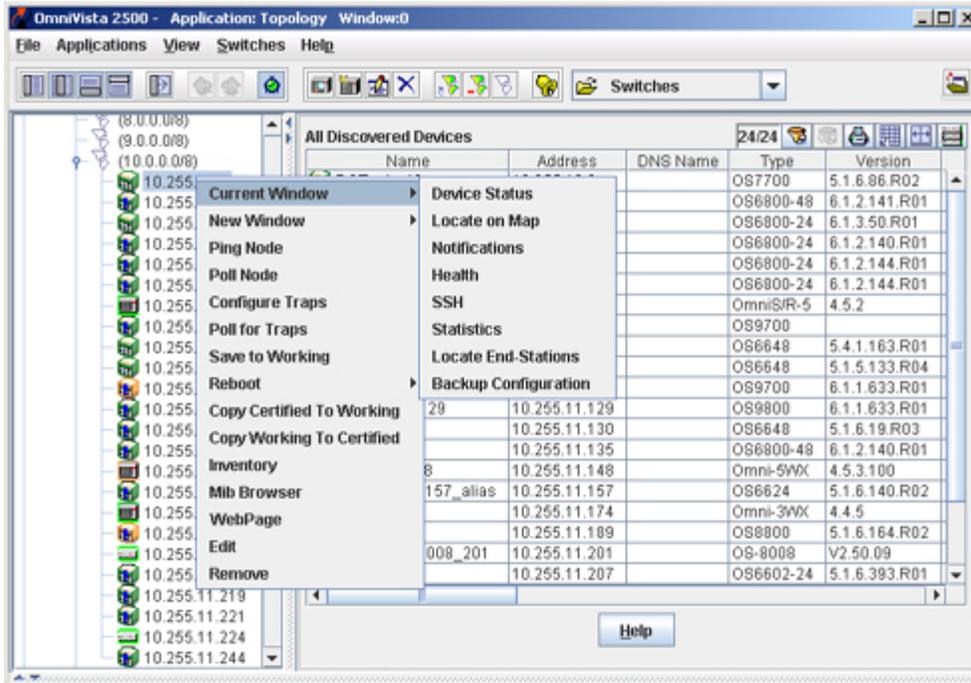
You can connect to a switch merely by selecting it in the Tree. When the connection is established, information about the switch displays, as shown below. Note that the information displayed is somewhat different for AOS devices, XOS devices or third-party devices.



Popup Menu in the Tree

You can click right on any device in the Tree to display a popup menu. Somewhat different versions of the Tree popup menu display for AOS devices, XOS devices, or third-party devices. The Tree popup menu for AOS devices is shown below. Note that all menu items on the Tree popup menu also appear on the popup menu in the list of All Discovered Devices (described above).

Tree Popup Menu for AOS Devices
Right-click on any device in the tree to display the menu.



Editing an Entry in the List of All Discovered Devices

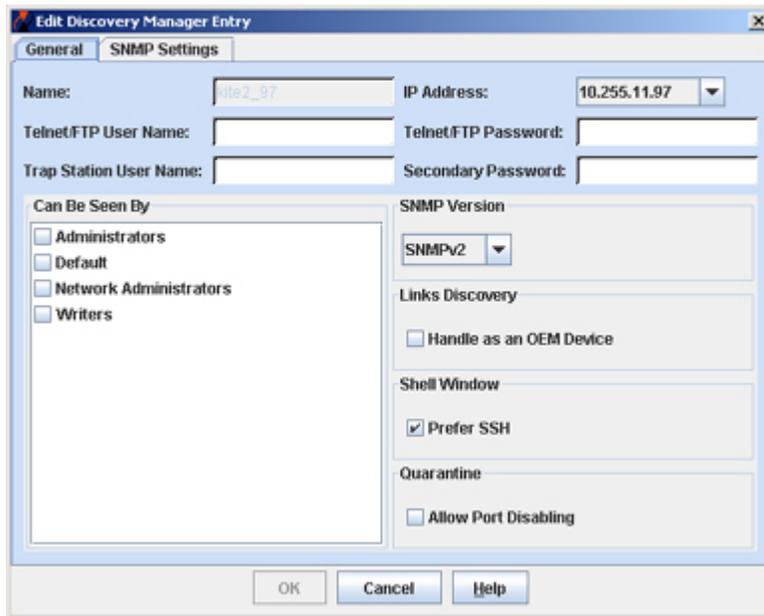
The Edit Discovery Manager Entry window, shown below, enables you to edit entries in the list of All Discovered Devices. You can redefine any field except the device name. Display the Edit Discovery Manager Entry window by

double clicking any single entry in the list of All Discovered Devices,

selecting one or more entries and clicking the edit icon  , or

selecting one or more entries, clicking right, and selecting **Edit** from the popup menu that displays.

Edit Discovery Manager Entry Window
Single Switch Selected



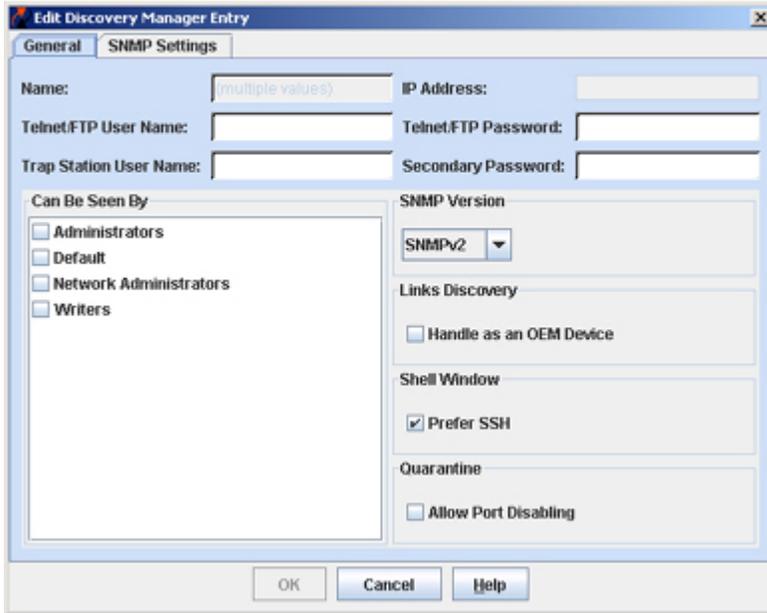
Editing Multiple Entries Simultaneously

It is possible to edit multiple entries in the list of All Discovered Devices simultaneously. To do this, select the devices in the list of All Discovered Devices and click the edit icon  , or click right and select **Edit** from the popup menu that displays. The Edit Discovery Manager Entry window displays "multiple values" in the **Name** field when more than one switch is selected. The "multiple values" message also displays in fields where the selected switches have different values, such as in the **SNMP Version** field shown below.

To edit all switches selected, merely enter values in the desired fields of the Edit Discovery Manager Entry window and click **OK**. The changes will apply to all switches selected. For example, if you were to enter **michael** in the **Telnet/FTP User Name** field and click **OK**, you would be specifying to OmniVista that the Telnet/FTP user name for all selected switches is **michael**. Note that any field you leave blank will retain

its former value. If you attempt to set a value that is not valid for all switches selected, an explanatory message will display and the change will not be made.

Edit Discovery Manager Entry Window Multiple Switches Selected



Why Edit an Entry?

You may want to edit entries in the list of All Discovered Devices for any or all of the following reasons:

To Redefine the Primary IP Address

When switches are autodiscovered 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 **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. The device's primary IP address will display as the device's address in the list of All Discovered Devices.

To Specify the Telnet and FTP User Name and Password

The **Telnet/FTP User Name** and **Telnet/FTP Password** fields enable you to specify the user name and password that OmniVista will use to establish FTP and Telnet sessions with the device. The user name and password specified will be used to auto-login to devices when Telnet sessions are established. They will also be used to perform FTP with the device when configuration files are saved and restored.

Firmware configuration files for XOS and 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. In order to FTP files, OmniVista must know the FTP login name and password that is defined on the switch. The **Telnet/FTP User Name** and **Password** fields enable you to specify this information to OmniVista.

Please Note:

If you do not define the Telnet/FTP login name and password, and you attempt to save, restore, or upgrade configuration files for XOS or 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.

If you do not define the Telnet/FTP login name and password, OmniVista will be unable to auto-login to the device when establishing Telnet sessions.

For OmniCore devices, the login name and password specified in these fields will be used to establish Telnet sessions and will be passed to the TrackView Element Manager automatically whenever TrackView is invoked.

To Define the Trap Station User Name (AOS Devices only)

This field enables you to specify the switch user name that will be used when an AOS device is configured to send traps to OmniVista. AOS devices require that a valid switch 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 nondefault 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 switch, the community string with which the switch is discovered is not guaranteed to be a valid switch user name, and thus is not guaranteed to be a valid switch user name for a trap station configuration entry. In this case, you should enter a valid switch user name in the **Trap Station User Name** field.

To Redefine Switch Access

The **Can Be Seen By** parameter specifies the OmniVista security group that has access to the device. The Edit Discovery Manager Entry window enables you to redefine the security group or to specify that all security groups have access.

To Redefine the SNMP Version

The Edit Discovery Manager Entry window enables you to redefine the SNMP version that OmniVista uses to communicate with AOS devices. XOS devices support SNMP version 1 only. AOS devices support SNMP version 1, SNMP version 2 or SNMP version 3.

To Specify How a Device's Links will be Discovered

The **Handle as an OEM Device** checkbox enables you to specify that you want a device's links to other data switches discovered automatically, 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. Such devices include the OmniPCX, OmniCore 5xxx switches, and third party devices.

Links to other switches are discovered automatically and displayed on Topology maps for all Alcatel devices that support adjacency protocols. AOS devices, XOS devices, and 61xx and 6300-24 devices all support adjacency protocols. In previous releases of OmniVista, devices that did not support adjacency protocols -- such as the OmniPCX, OmniCore 5xxx switches, and third party devices -- were discovered and displayed on Topology maps, but links from these devices to other switches had to be added manually.

As stated, the **Handle as an OEM Device** checkbox now enables you to use the new "endstation search" functionality from the Locator application to automatically discover links for such devices. When the **Handle as an OEM Device** checkbox is enabled, and the 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 displayed on Topology maps automatically.

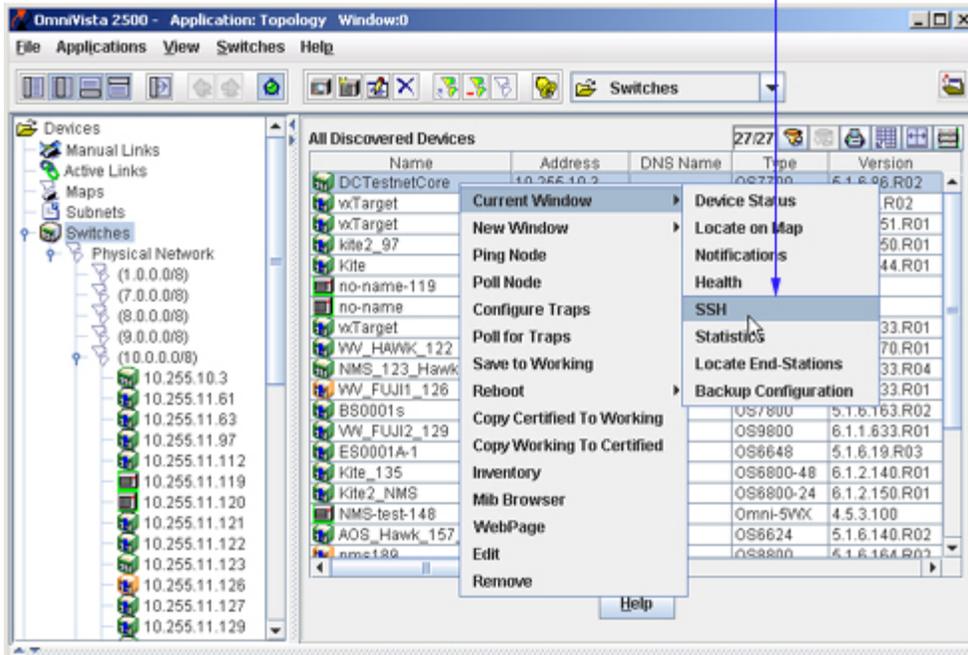
Note: This approach works well for switches located at the edge of the network that do not support adjacency protocols. However, when a series of such switches are interconnected at the core of a network, this approach may "discover" more links than are meaningful. As an example, consider a series of such switches 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 switches 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.

To Specify SSH as the Default Command Line Interface

OmniVista's Telnet 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 **Prefer SSH** checkbox is enabled, 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 the **Prefer SSH** checkbox is not enabled, 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. OmniVista popup menus, such as the one shown below, will automatically display the default command line interface for the device: **Telnet** or **SSH**. When selected, the Telnet application will open and a connection of the configured type will be established automatically.

Note: 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. XOS devices, OmniCore devices, and OmniStack 6124/6148 and 6300-24 devices are not SSH-capable.

When the Prefer SSH checkbox is enabled, OmniVista popup menus will display an SSH option instead of the default Telnet option.



To Specify Port Disabling on a Device

By default, all switches allow port disabling. However, if you want to enable port disabling for a specific device using OmniVista, click the Allow Port Disabling checkbox.

To Specify the Correct Write Community Name (SNMP Settings Tab)

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

In like manner, if someone changes a switch's read community name or write community name after the switch has been autodiscovered, use the Edit Discovery Manager Entry window to redefine the community name to OmniVista. Note that OmniVista will lose connection with a switch if its read community name is changed; when the correct read community name is specified to OmniVista the connection will be automatically reestablished.

To Redefine SNMP Parameters (SNMP Settings Tab)

The SNMP Settings tab of the Edit Discovery Manager Entry window enables you to redefine SNMP parameters in addition to the write community name. You can redefine parameters for SNMP version 1, SNMP version 2, and SNMP version 3. Click here for further information.

Editable Fields

To edit entries, redefine the desired fields and click the **OK** button. Note that you cannot change the device name.

The General Tab

IP Address field

Set this combo box to the IP address that you want OmniVista to use as the device's primary IP address. The IP address combo box displays all IP addresses associated with the device that responded to OmniVista's ping during autodiscovery. The device's primary IP address will display as the device's address in the list of All Discovered Devices.

Telnet/FTP User Name and Telnet/FTP Password Fields

Enter the switch's Telnet/FTP login name in the **Telnet/FTP User Name** field and enter the switch's Telnet/FTP password in the **Telnet/FTP Password** field. Note that a more complete discussion of these fields is found above.

Please Note: These fields enable you to inform OmniVista of the switch's Telnet/FTP user name and password. A switch's Telnet/FTP user name and password cannot be configured from OmniVista. The Telnet/FTP user name and password must be configured directly on the switch.

Trap Station User Name field (AOS devices only)

The **Trap Station User Name** field enables you to specify the switch user name that will be used when an AOS device is configured to send traps to OmniVista. AOS devices require that a valid switch user name be specified with the trap station configuration entry. Note that a more complete discussion of this field is found above.

Can Be Seen By field

The **Can Be Seen By** field specifies the security permissions that are required for viewing the switch in the list of all discovered devices. OmniVista is shipped with predefined user groups that have various levels of security permissions. The network administrator may have modified these groups or created new ones. (The Security application *Users and Groups* enables you to view and configure security permissions for users.) Checkboxes for all existing user groups are displayed. Click the checkbox by each user group that you want to have access to the switch. Alternatively, if you do not click any checkbox, the switches will be viewable by everyone. The predefined user groups are as follows:

Everyone. Everyone that logs into OmniVista will be able to view the switch in the list of discovered devices.

Network Administrators. Only users that have administrative permissions will be able to view the switch in the list of discovered devices.

Writers. Users that have read/write permissions will be able to view the switch in the list of discovered devices. Note that users with administrative permissions also have read/write permissions and thus will also be able to view the switch in the list of discovered devices.

Default. Users that have default permission (the default permission is read) will be able to view the switch in the list of discovered devices. Note that users with administrative permissions also have read permission and thus will also be able to view the switch in the list of discovered devices.

SNMP Version

The SNMP Version combo box displays the SNMP version that OmniVista is using to communicate with the switch. For XOS devices, which support SNMP version 1 only, the combo box is always set to **SNMPv1** (SNMP version 1) and cannot be changed. For AOS devices, the combo box defaults to SNMP version 2, but can be changed to SNMP version 1 or SNMP version 3. (AOS devices support SNMP version 1, SNMP version 2, or SNMP version 3.) To change the SNMP version that OmniVista uses to communicate with an AOS device, merely set the combo box to the desired SNMP version and click **OK**.

Handle as an OEM Device checkbox

The **Handle as an OEM Device** checkbox, when enabled, specifies that you want a device's links to other data switches discovered automatically, 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. Such devices include the OmniPCX, OmniCore 5xxx switches, and third party devices. Note that a more complete discussion of this field is found above.

Prefer SSH checkbox

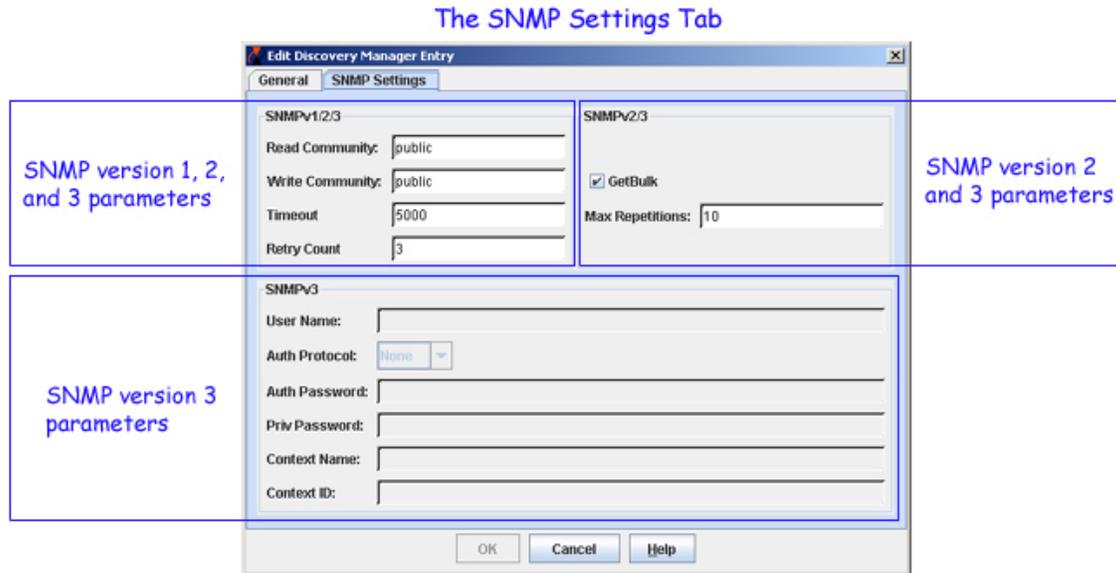
The **Prefer SSH** checkbox, when enabled, specifies that SSH (Secure Shell) will be used as the default command line interface for the device, and that **SSH** will display on OmniVista popup menus instead of **Telnet**. In addition, Secure Shell FTP will be used as the default FTP method in Resource Manager. If the **Prefer SSH** checkbox is not enabled, Telnet will be used as the default command line interface for the device, and **Telnet** will display on OmniVista popup menus. Regular FTP will be used as the default FTP method in Resource Manager. OmniVista's Telnet 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. Note that a more complete discussion of this field is found above.

Allow Port Disabling checkbox

The **Allow Port Disabling** checkbox, when enabled, specifies that port disabling is allowed for the device using OmniVista. If the **Allow Port Disabling** checkbox is not enabled, you cannot disable ports for the device using OmniVista.

The SNMP Settings Tab

The SNMP Settings tab displays parameters for SNMP versions 1, 2, and 3. These parameters are cumulative in that SNMP version 1 supports only version 1 parameters, SNMP version 2 supports version 1 and version 2 parameters, and SNMP version 3 supports version 1, version 2, and version 3 parameters. Only those parameters that are supported by the current version of SNMP will be active in the SNMP Settings tab. Each parameter is explained below.



SNMP Versions 1, 2, and 3 Parameters

Read Community and Write Community

In the **Read Community** field, enter the switch's get community name. The get community name enables you to read information from the switch. In the **Write Community** field, enter the switch's set community name. The set community name enables you to write information to the switch. If the switch'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.)

Please Note:

Get and set community names are not configurable from OmniVista. Get and set community names can only be configured by logging onto the switch.

When you use SNMP Version 3, get and set community names are ignored.

Timeout

The **Timeout** field specifies 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.

Retry Count

The **Retry Count** field specifies the number of times that OmniVista will attempt to connect to a switch.

SNMP Versions 2 and 3 Parameters

GetBulk checkbox

The **GetBulk** checkbox is enabled by default. You can disable Get Bulk operations by clicking the checkbox to uncheck it. 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. 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 value in the **Max Repetitions** field determines the number of rows of table data that the Get Bulk operation will request in each Get Next operation.

SNMP Version 3 Parameters

User Name

Enter the SNMP version 3 user name in this field.

AuthProtocol

Set this field to **None**, **MD5**, or **SHA** to specify the authentication protocol OmniVista will use for SNMP communication with the switch. **MD5** (or HMAC-MD5-96) and **SHA** (or HMAC-SHA-96) are the two authentication protocols that have been defined for SNMP version 3.

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. The MD5 and SHA authentication protocols produce the MAC in a similar, but not an identical, manner.

Note that the **Auth Password** and **Priv Password** fields activate when the authentication protocol is set to something other than **None**. The Privacy Password field activates because privacy can only be used when authentication is also used. The Authentication password field activates because the authentication password is used as the "secret key" mentioned above. For MD5 the secret key should be 16 octets; for SHA the secret key should be 20 octets. Note that this implies that stronger authentication is provided by the SHA protocol, and SHA should be used instead of MD5 when possible.

Auth Password

Enter the password (in hex) 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.

Priv Password

SNMP version 3 uses the CBC-DES Symmetric Encryption Protocol for privacy. Enter the password in the **Priv Password** field (in hex) 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.

Important Note: The switch uses a single password as both the **Auth Password** and the **Priv Password**. This means that the same password should be entered in these two fields. You can identify the password to enter by using the switch CLI command **configuration snapshot aaa**. This command will show the "authkey" for each switch user. The authkey is a hex value computed from the user's password. (The user's password is established with the CLI command **user**.) If you want both authentication and privacy encryption, enter the authkey in both the **Auth Password** and the **Priv Password** fields.

Context Name

Enter 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

Enter 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.

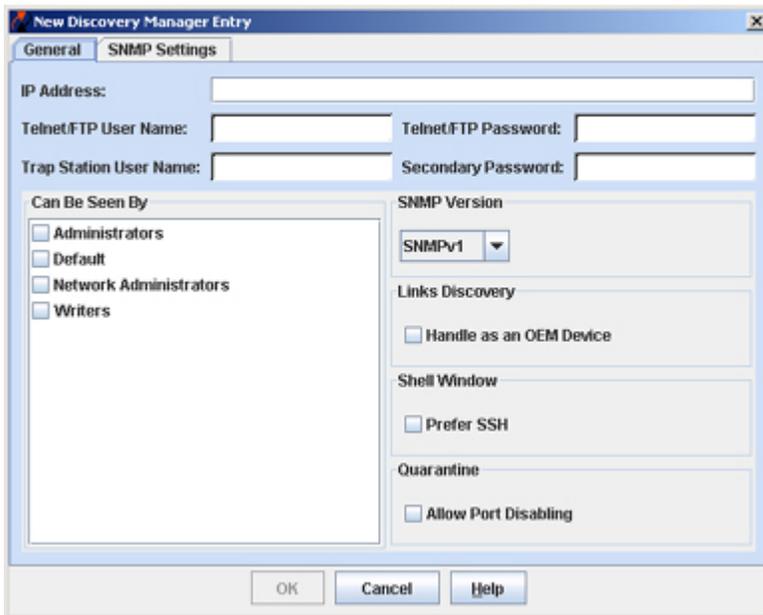
Important Note: Neither the **Context Name** nor the **Context ID** are required for AOS, XOS, or default third-party devices supported by OmniVista. Leave these fields blank unless you are using a non-default third-party device that requires definition of a Context Name and Context ID.

Adding a New Entry

The New Discovery Manager Entry window, shown below, enables you to manually add a new switch to the list of All Discovered Devices. Display the New Discovery Manager Entry window by clicking the "Create a New Switch Entry" icon  or by selecting **New** on the Switches menu. Follow the steps below to add a new switch to the list of All Discovered Devices.

Note: The following fields on the New Discovery Manager Entry window cannot be configured until OmniVista has connected to the switch: the **Trap Station User Name** field and the **SNMP Version** field on the General tab, and any SNMP parameters used exclusively by SNMPv2 or SNMPv3 on the SNMP Setting tab. To configure these fields, edit the switch entry after OmniVista has connected to the switch.

New Discovery Manager Entry Window
General Tab



1. In the **IP Address** field, enter the IP address of the new switch.
2. Enter the switch's Telnet/FTP login name in the **Telnet/FTP User Name** field and the switch's Telnet/FTP password in the **Telnet/FTP Password** field. These fields specify the user name and password that OmniVista will use to establish FTP and Telnet sessions with the device. The user name and password specified will be used to auto-login to devices when Telnet sessions are established. They will also be used to perform FTP with the device when configuration files are saved and restored. Click here for a more complete discussion of these fields.

Please Note: These fields enable you to inform OmniVista of the switch's Telnet/FTP user name and password. A switch's Telnet/FTP user name and password cannot be configured from OmniVista. The Telnet/FTP user name and password must be configured directly on the switch.

3. Set the **Can Be Seen By** field to specify the security permissions that will be required for viewing the switch in the list of All Discovered Devices AFTER it is added. OmniVista is shipped with predefined user groups that have various levels of security permissions. The network administrator may have modified these groups or created new ones. (The Security application *Users and Groups* enables you to view and configure security permissions for users.) Checkboxes for all existing user groups are displayed. Click the checkbox by each user group that you want to have access to the switch. Alternatively, if you do not click any checkbox, the switches will be viewable by everyone. The predefined user groups are as follows:

Everyone. Everyone that logs into OmniVista will be able to view the switch in the list of discovered devices.

Network Administrators. Only users that have administrative permissions will be able to view the switch in the list of discovered devices.

Writers. Users that have read/write permissions will be able to view the switch in the list of discovered devices. Note that users with administrative permissions also have read/write permissions and thus will also be able to view the switch in the list of discovered devices.

Default. Users that have default permission (the default permission is read) will be able to view the switch in the list of discovered devices. Note that users with administrative permissions also have read permission and thus will also be able to view the switch in the list of discovered devices.

4. You can select the SNMP version that OmniVista uses to communicate with the switch from the the **SNMP Version** combo box. For XOS devices, which support SNMP version 1 only, the combo box is always set to **SNMPv1** (SNMP version 1) and cannot be changed. For AOS devices, the combo box defaults to SNMP version 2, but can be changed to SNMP version 1 or SNMP version 3. (AOS devices support SNMP version 1, SNMP version 2, or SNMP version 3.) To change the SNMP version that OmniVista uses to communicate with an AOS device, merely set the combo box to the desired SNMP version and click **OK**.

5. You can enable the **Handle as an OEM Device** checkbox, if you want a device's links to other data switches discovered automatically, 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. Such devices include the OmniPCX, OmniCore 5xxx switches, and third party devices.

6. You can enable the **Prefer SSH** checkbox, if you want SSH (Secure Shell) to be used as the default command line interface for the device, which will be displayed as a popup menu item instead of **Telnet**. Secure Shell FTP will be used as the default FTP method in Resource Manager. If the **Prefer SSH** checkbox is not enabled, Telnet will be used as the default command line interface for the device, and **Telnet** will display on OmniVista popup menus. Regular FTP will be used as the default FTP method in Resource Manager. OmniVista's Telnet 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. Note that a more complete discussion of this field is found above.

7. You can enable the **Allow Port Disabling** checkbox to allow port disabling for the device using OmniVista. If the **Allow Port Disabling** checkbox is not enabled, you cannot disable ports for the device using OmniVista.

8. Click the **SNMP Settings** tab to display the SNMP Settings page of the New Discovery Manager Entry window, shown below.

New Discovery Manager Entry Window SNMP Settings Tab

The screenshot shows the 'New Discovery Manager Entry' window with the 'SNMP Settings' tab selected. The window is divided into three main sections for configuring different SNMP versions. The 'SNMPv1/2/3' section includes fields for 'Read Community', 'Write Community', 'Timeout' (5000), and 'Retry Count' (3). The 'SNMPv2/3' section has a checked 'GetBulk' checkbox and a 'Max Repetitions' field (0). The 'SNMPv3' section includes fields for 'User Name', 'Auth Protocol' (None), 'Auth Password', 'Priv Password', 'Context Name', and 'Context ID'. The window has 'OK', 'Cancel', and 'Help' buttons at the bottom.

9. In the **Read Community** field, enter the new switch's get community name. The get community name enables you to read information from the switch. In the **Write Community** field, enter the new switch's set community name. The set community name enables you to write information to the switch. If the switch'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.)

Please Note: These fields enable you to inform OmniVista of the switch's SNMP get and set community names. A switch's get and set community names cannot be configured from OmniVista. They must be configured directly on the switch.

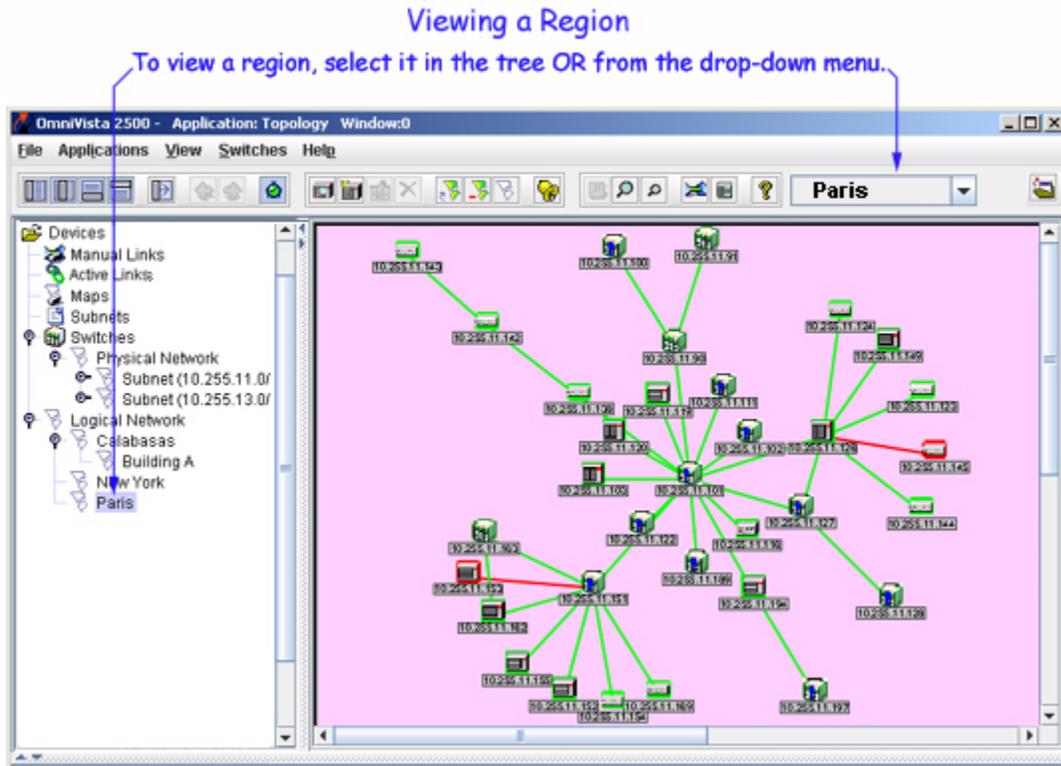
10. In the **Timeout** field, specify the time period, in milliseconds, that OmniVista will wait for the switch to respond to a connection request before assuming that the request has timed-out.

11. In the **Retry Count** field, specify the number of times that OmniVista will attempt to connect to the switch.

12. Click the **OK** button. The new switch is added to the list of All Discovered Devices and to the Tree.

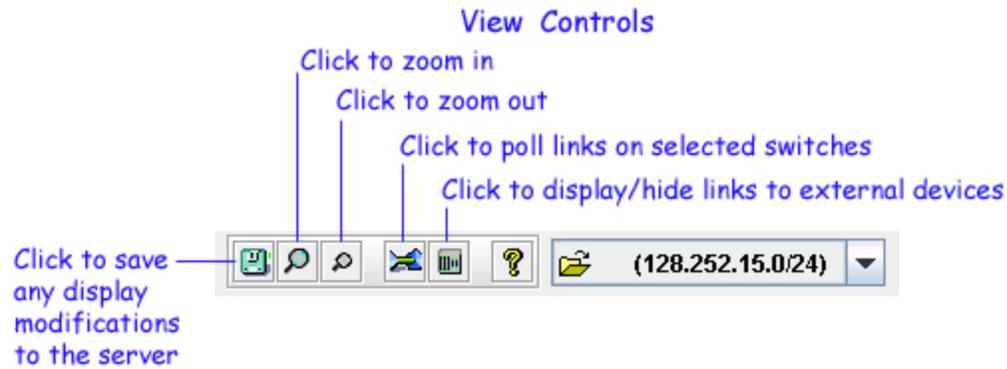
Viewing the Network

The Topology application enables you to display a map of any network region, including the overall Physical Network, the overall Logical Network, or any individual subnet or region therein. Color coding in regional maps provides status information on each region, device, and link displayed. Specific information about the links in each region can be viewed. Popup menus provide further functionality. To view a map of any network region, select the region in the Tree or in the combo box shown below. When selected, the regional map displays with the background color and background image specified when the region was created or edited.



View Controls

As shown below, controls at the top of the viewing window enable you to zoom in and out of regional maps. You can toggle display of any external links that exist within a region. ("External" links are links to devices that are not part of the region displayed.) You can poll links on selected switches to gather current information about the links. And the Save icon enables you to save any changes you make to the display. (If you make changes to a regional map and do not save them, you will be asked if you want to validate your modifications when you exit the display.)



Note: Clicking the "Poll Links on Selected Switches" icon causes an immediate poll of all links associated with each selected device in the regional map, and all devices connected to them. If no devices are selected, you will be asked if you wish to poll the links for all devices in the map. This feature is useful when you want to quickly refresh the link data for the selected devices or for the entire map. The success or failure of the poll is reported in the Status Panel for each individual device.

Viewing Link Information

While viewing a regional map, you can display information about the links in the region. To do this, merely place the cursor on the connection you wish to view. A table listing the individual links in the connection displays, as shown below. The fields in this table are explained below.

Note that connections composed of more than one link are represented by thicker lines in the display. These are termed aggregate connections. For example, in the screen below the connection between device 10.255.11.156 and device 10.255.11.155 displays with a thicker line because this connection consists of more than one link. Compare this with the thinner line used to display other connections, which consist of a single link.

Place the cursor on a connection you want to view. A table listing all links in the connection displays.

Links for 10.255.11.111 - 10.255.11.112

| Origin | IP | Slot/Port | LAG Id | IP | Slot/Port | LAG Id | Type | Status | VLAN Id |
|--------|---------------|-----------|--------|---------------|-----------|--------|----------------|--------|---------|
| AMAP | 10.255.11.112 | 1/22 | | 10.255.11.111 | 1/11 | | ethernetCsmacd | Up | 70 |

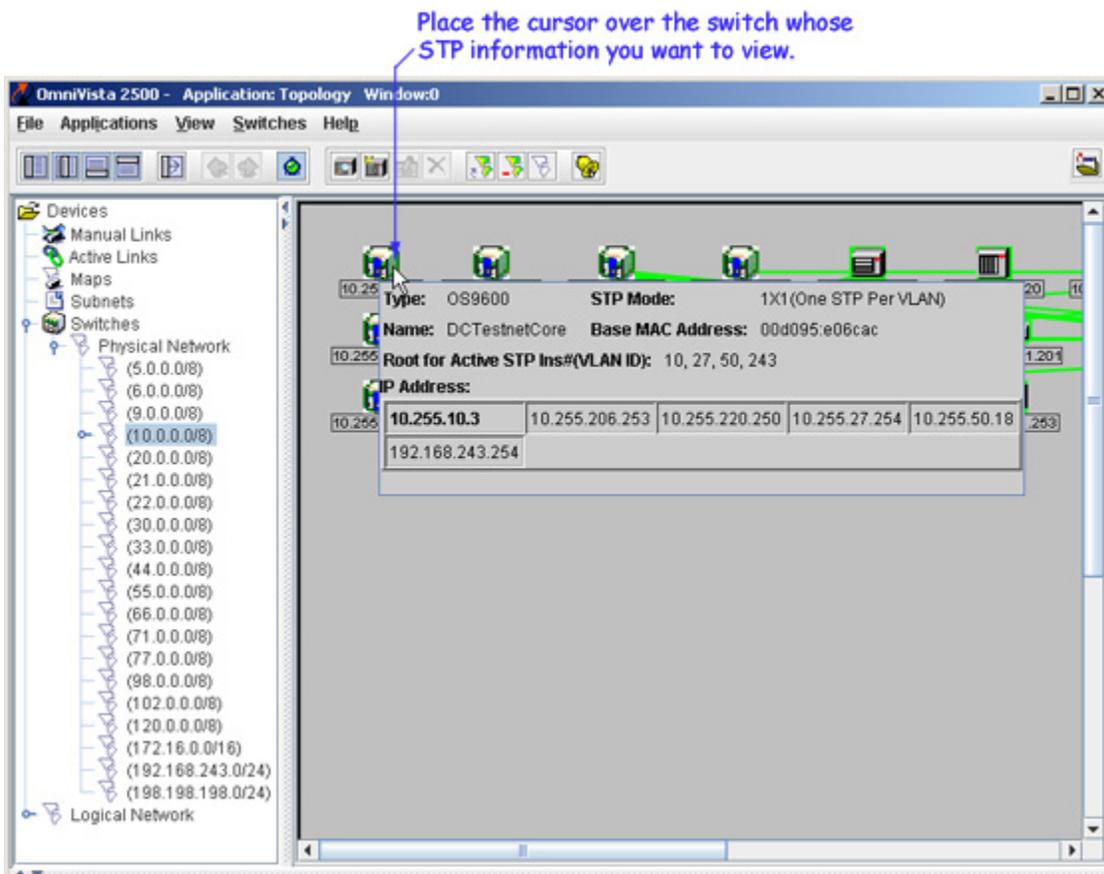
The fields in the Link Information table are explained below.

| The Link Origin, Switch IP address, slot/port, and link aggregation ID (if any) for the first link in the connection. | | Switch IP address, slot/port, and link aggregation ID (if any) for the second link in the connection. | | | The type of connection, its status, and VLAN Id. | | | | |
|---|---------------|---|--------|---------------|--|--------|----------------|--------|---------|
| Links for 10.255.11.111 - 10.255.11.112 | | | | | | | | | |
| Origin | IP | Slot/Port | LAG Id | IP | Slot/Port | LAG Id | Type | Status | VLAN Id |
| AMAP | 10.255.11.112 | 1/22 | | 10.255.11.111 | 1/11 | | ethernetCsmacd | Up | 70 |

Note: Links to OEM devices will always display "1/1" on the OEM side in the **Slot/Port** column regardless of what slot and port on the OEM device is actually used.

Viewing STP Information

While viewing a regional map, you can display STP information for a switch. To view this information, move the mouse over a switch in the OmniVista's Topology Map graphical display. This information is only displayed if STP information has already been collected from the switch. Detailed STP Port information is available by right-clicking on a switch.



Note: The tooltip with the STP information will be displayed for AOS and XOS devices only.

Color Coding in Regional Maps

Whenever a regional map is displayed, color coding provides status information about the devices, links, or subregions displayed. Generally, color coding indicates the following:

Green indicates a device, link, or region is up and therefore in the "normal" state.

Red indicates a device, link, or region is down and therefore in the "critical" state.

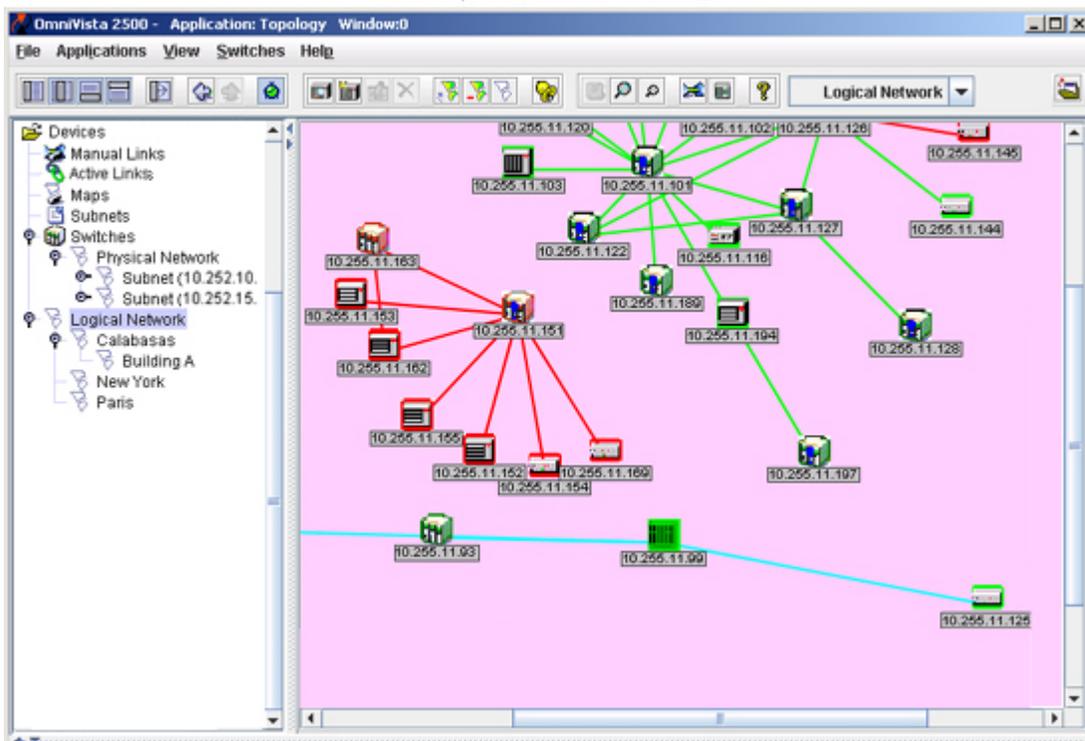
Blue indicates a link is in an "unknown" state. When OmniVista receives a trap indicating that a the switch is rebooting, the link state changes to "unknown", and the link is displayed in blue until the next poll, or until a trap is received that shows the link state.

Orange indicates a device, an aggregate link, or a region is in the "warning" state.

Note: For proper display of links in AOS switches, linkUp and linkDown traps must be enabled for each individual port.

Color coding is explained in greater detail below.

Example of Color Coding



Individual Link Display Colors — — —

Individual links can display green (to indicate they are up), red (to indicate they are down), or blue (to indicate their status is unknown).

Aggregate Connection Display Colors

Aggregate connections, which are connections composed of more than one link, can display green (if every individual link in the connection is up), red (if every individual link in the connection is down), blue (to indicate the status of every individual link in the connection is unknown), or orange (to indicate all other cases -- for example, if one end of a link is up and the other end is down).

Note: Discovered links to OEM devices will always be displayed in blue.

Device Display Colors

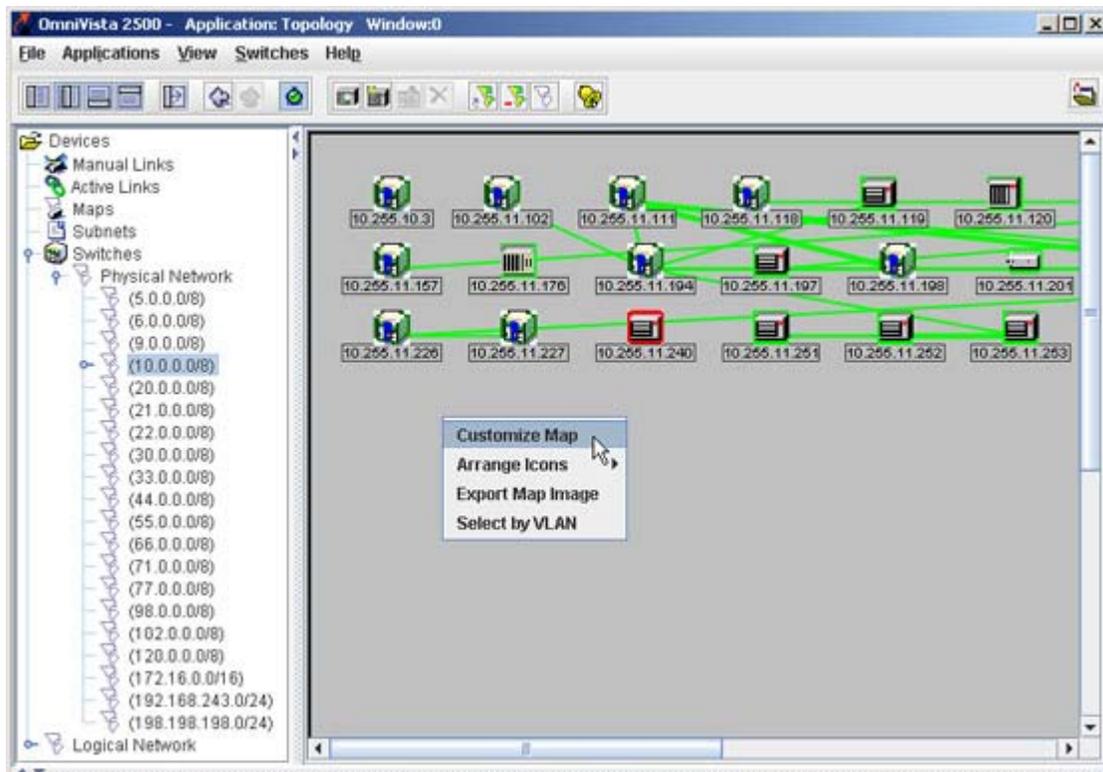
A device can display green (to indicate it is up and in the normal state), orange (to indicate the device has sent at least one warning or critical trap and is thus in the warning state), or red (to indicate the device is down and in the critical state). Note that device status is also reported in the list of All Discovered Devices (in the Status column).

Region Display Colors

Regions can display green (if every device and link in the region is up and in the normal state), orange (if at least one device or link in the region is in a warning or unknown state), or red (if at least one device or link in the region is down and in the critical state). Note that region status is also reported in the Maps List (which you can view by selecting **Maps** in the Tree).

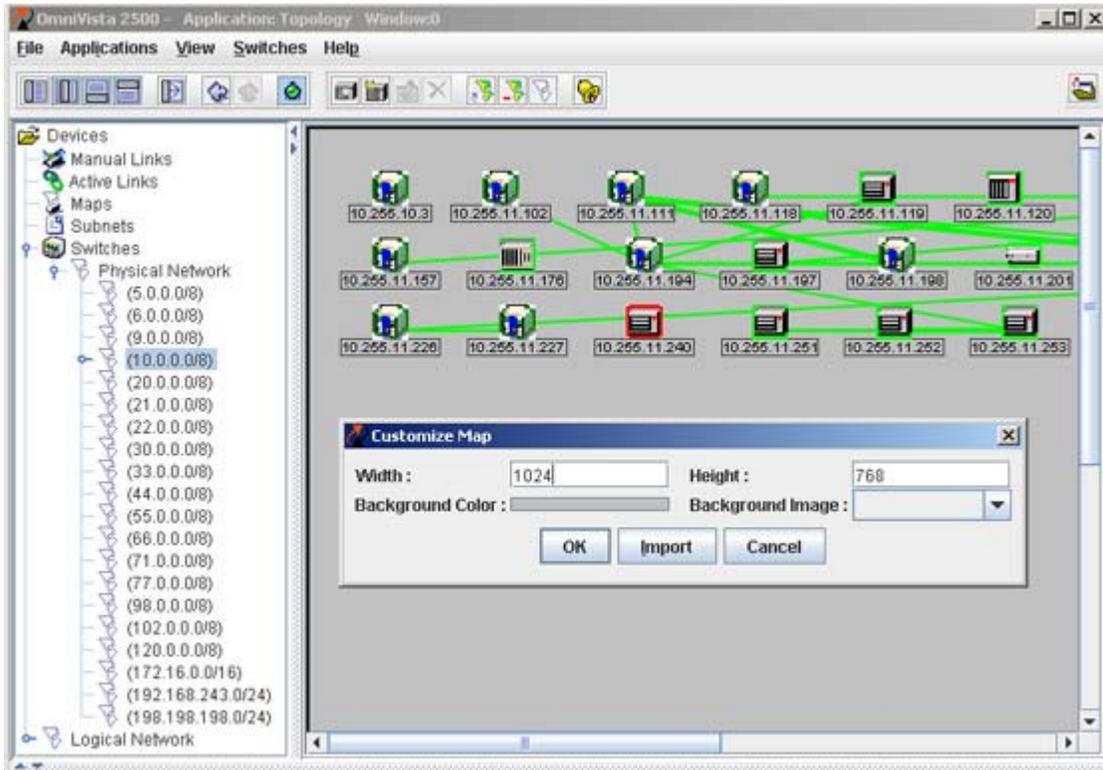
Customizing Regional Maps

When a regional map is displayed, you can customize its appearance by right-clicking anywhere in the background of the map. A pop-up menu is displayed (as shown below). Each menu item is described below.



The Customize Map Menu Item

This menu item enables you to customize the appearance of a regional map. When you click **Customize Map**, the Customize Map dialog box is displayed. This dialog box enables you to customize the size of the viewing window, the background color of the regional map, and the background image against which the map is displayed.



Make the required changes in the following fields, and then click the **OK** button. The customized map will be displayed.

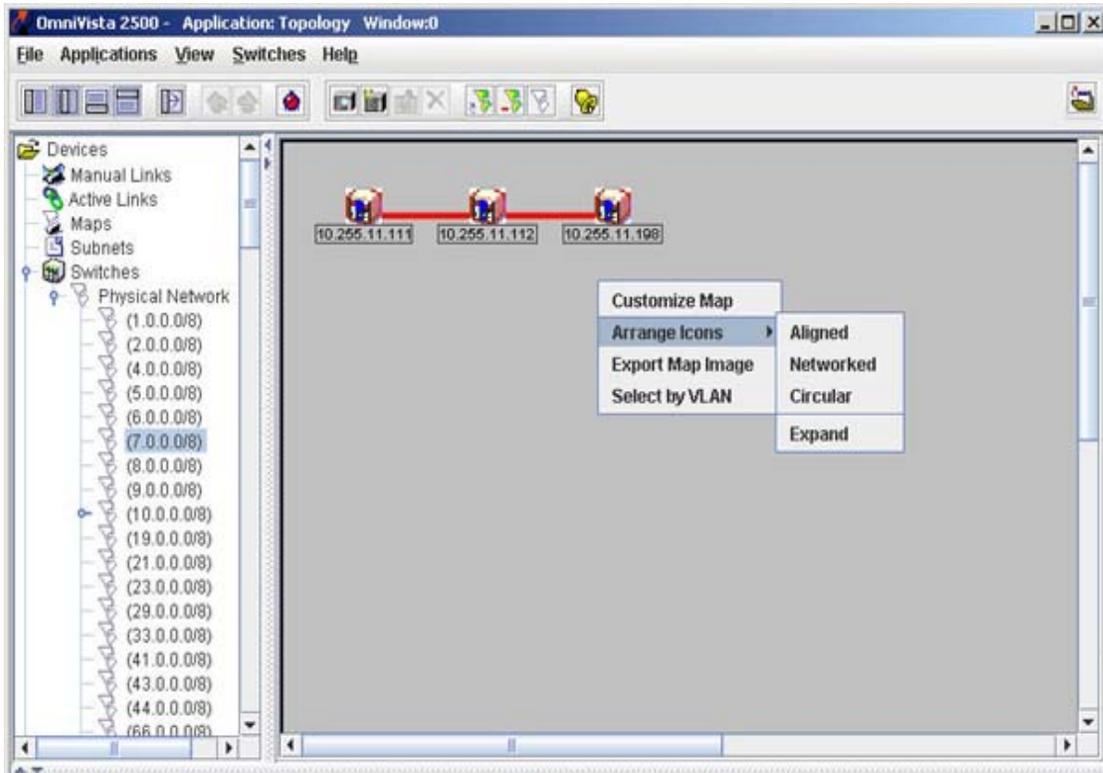
Width and **Height** fields. Redefine the size of the viewing window that displays the regional map by entering the desired width (in pixels) in the **Width** field and the desired height (in pixels) in the **Height** field.

Background Color field. Redefine the background color of the regional map. To do this, click left anywhere in the color displayed by the **Background Color** field. The Color Chooser displays. The Color Chooser enables you to define a background color by selecting a color swatch from a group of pre-defined swatches, or by using the HSB (hue saturation brightness) color model, or by using the RGB (red green blue) color model.

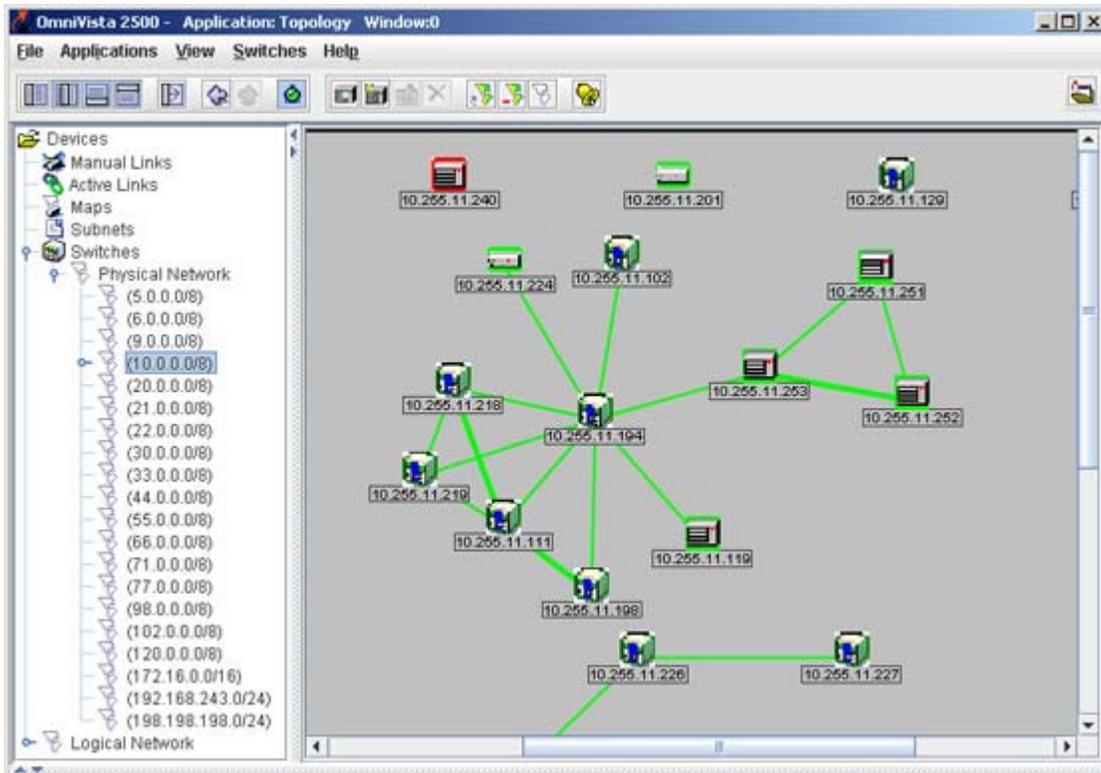
Background Image field. Add or redefine the background image for the regional map. To do this, set the **Background Image** combo box to the desired image. Note that all background images must be imported into OmniVista before they can be used. You can import background images by clicking the **Import** button on the Customize Map window. Refer to Importing Background Images for Regions for more information.

The Arrange Icons Menu Item

This menu item enables you to automatically arrange the device icons displayed in a regional map.

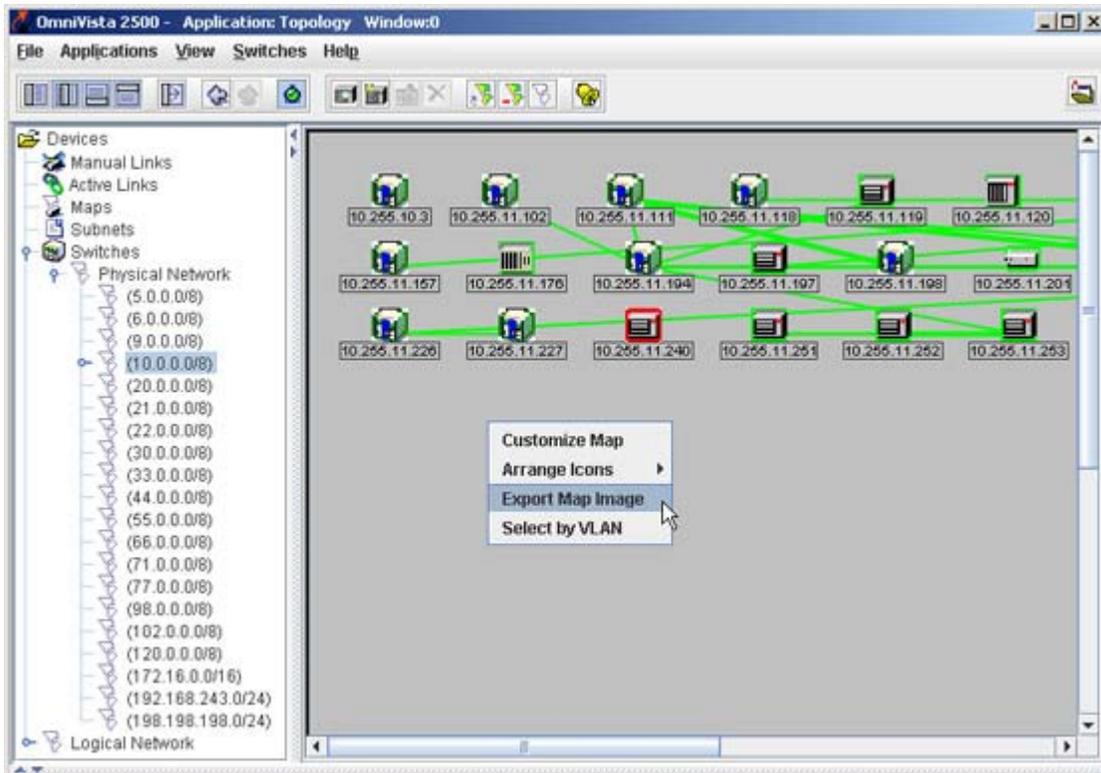


Arrange Icons > Aligned arranges icons in rows with no overlap (as shown below). **Arrange Icons > Networked** centers the icon that has the most connections and arranges other icons according to their connections (as shown above). **Arrange Icons > Circular** arranges icons in as much of a circle as possible. The **Arrange Icons > Expand** menu item expands the space between icons.

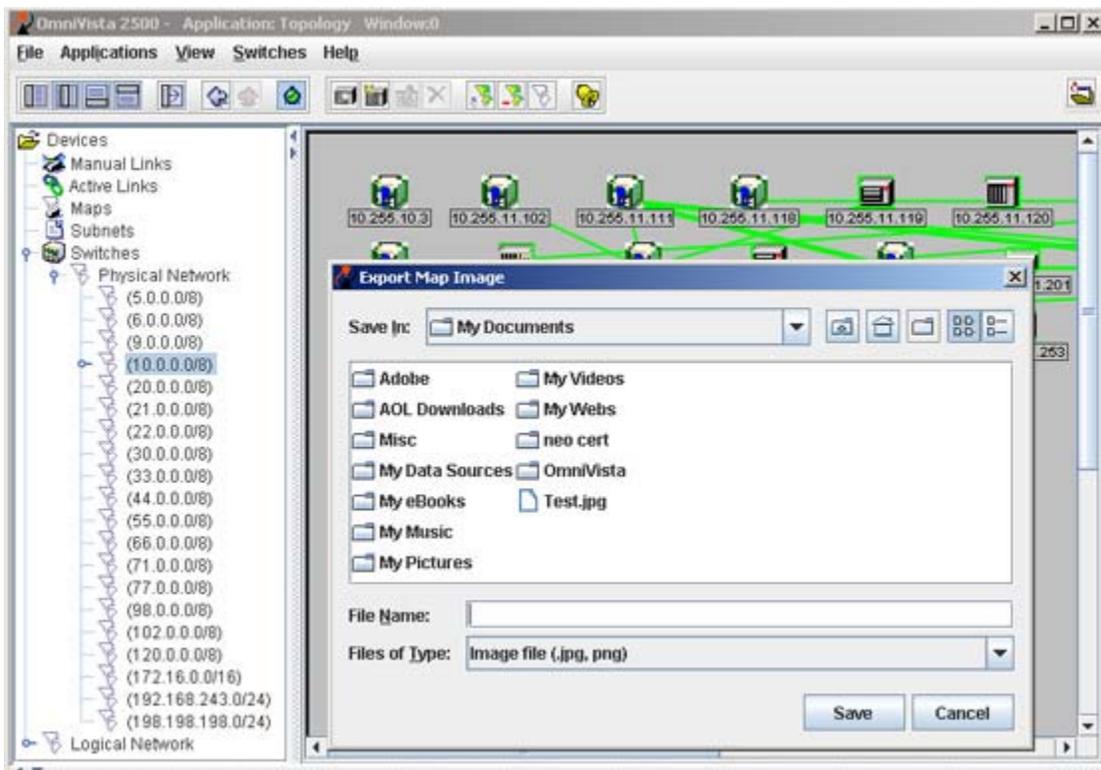


The Export Map Image Menu Item

This menu item enables you to export the map to the desired location.

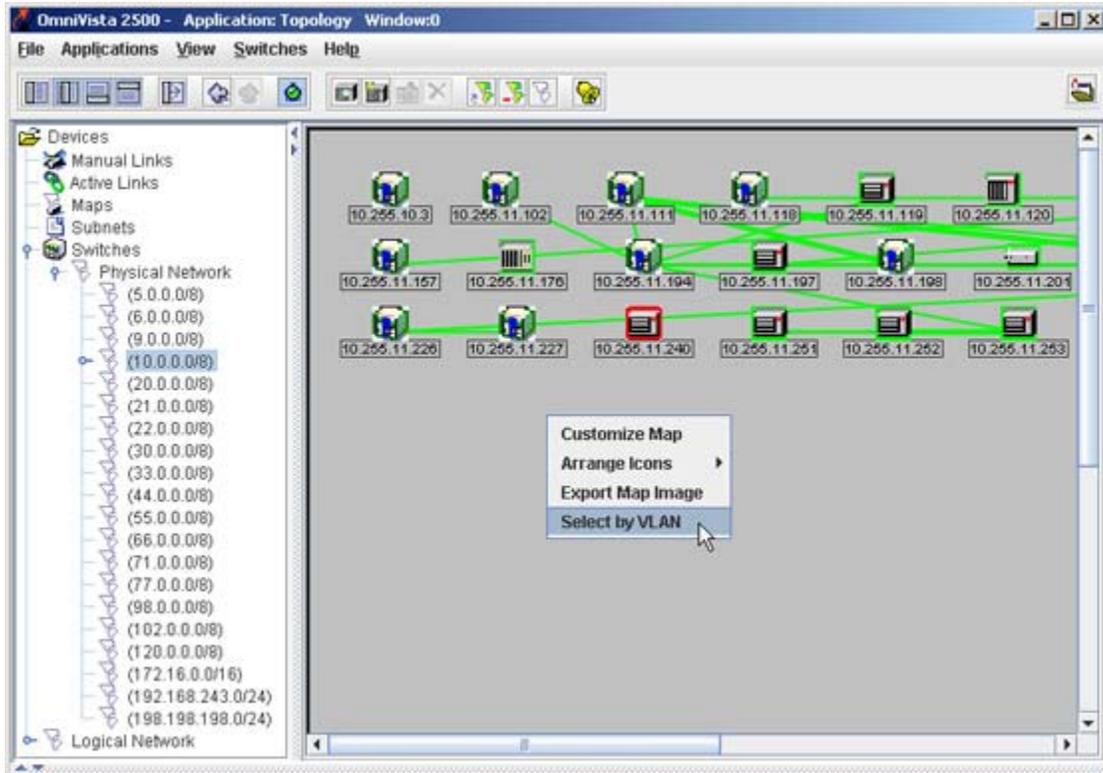


When you click **Export Map Image**, the Export Map Image dialog box is displayed (as shown below). Save the map in the required location.

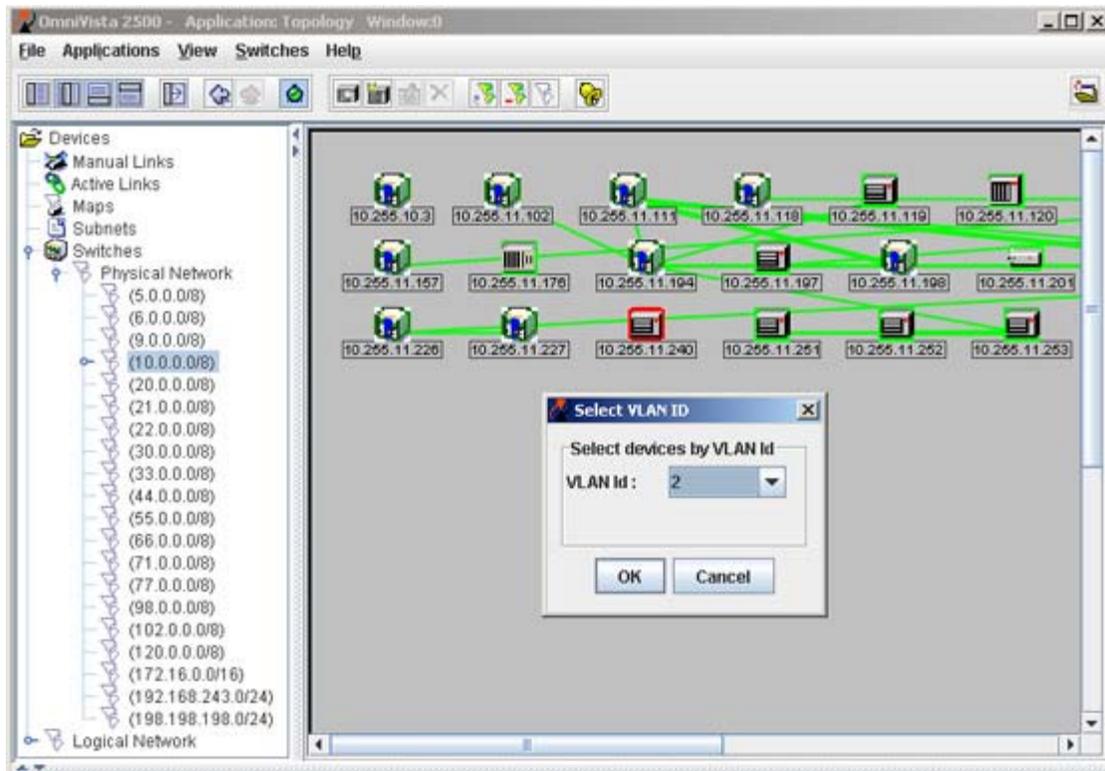


The Select by VLAN Menu Item

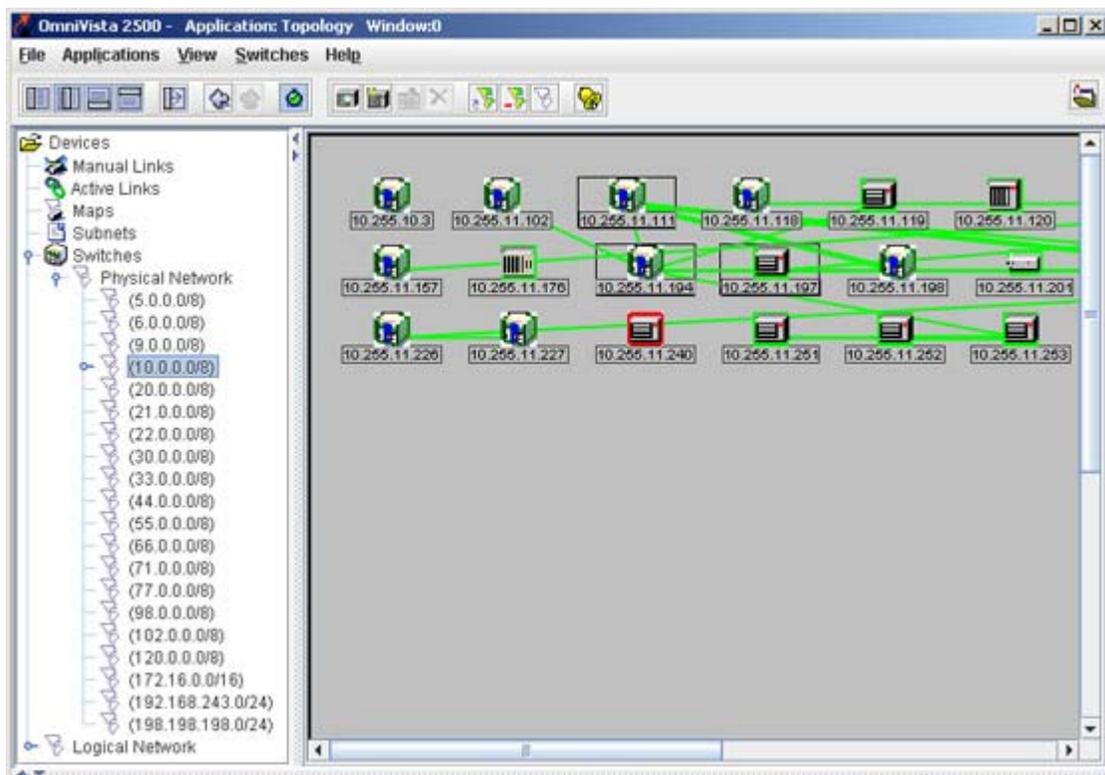
This menu item enables you to create a map from an existing map that contains only switches that belong to a given VLAN.



When you click **Select by VLAN**, the Select VLAN ID dialog box is displayed.



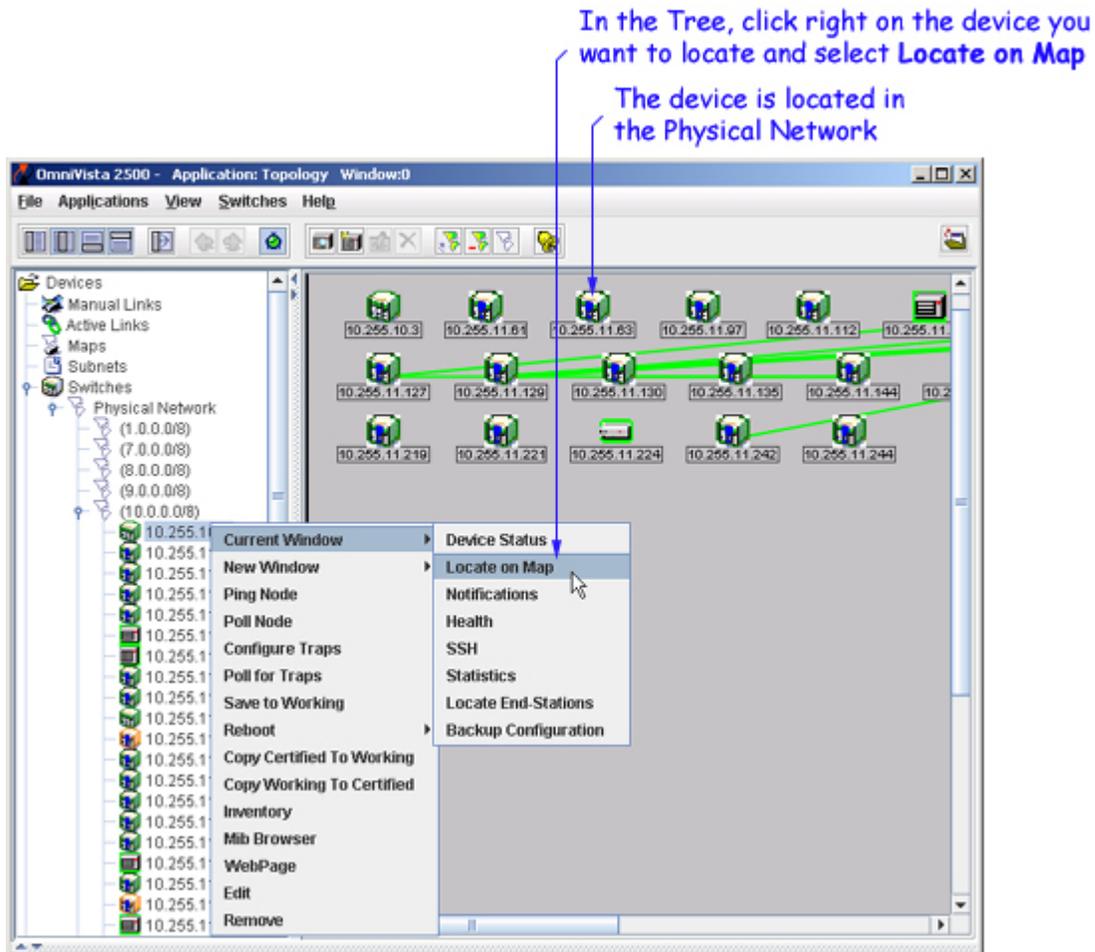
Select the required VLAN ID from the **VLAN ID** drop-down list, and click **OK**. The switches that belong to the selected VLAN will be highlighted (as shown below).



Locating a Specific Device

You can locate a specific device in the Physical Network by clicking right on the device in the Tree and then selecting **Locate on Map**, either in the **Current Window** or in a **New Window**. The device will be located in the Physical Network display, selected, and centered in the viewing window, as shown below. Note that the **Locate on Map** menu item is available on pop-up menus throughout OmniVista.

Locating a Specific Device

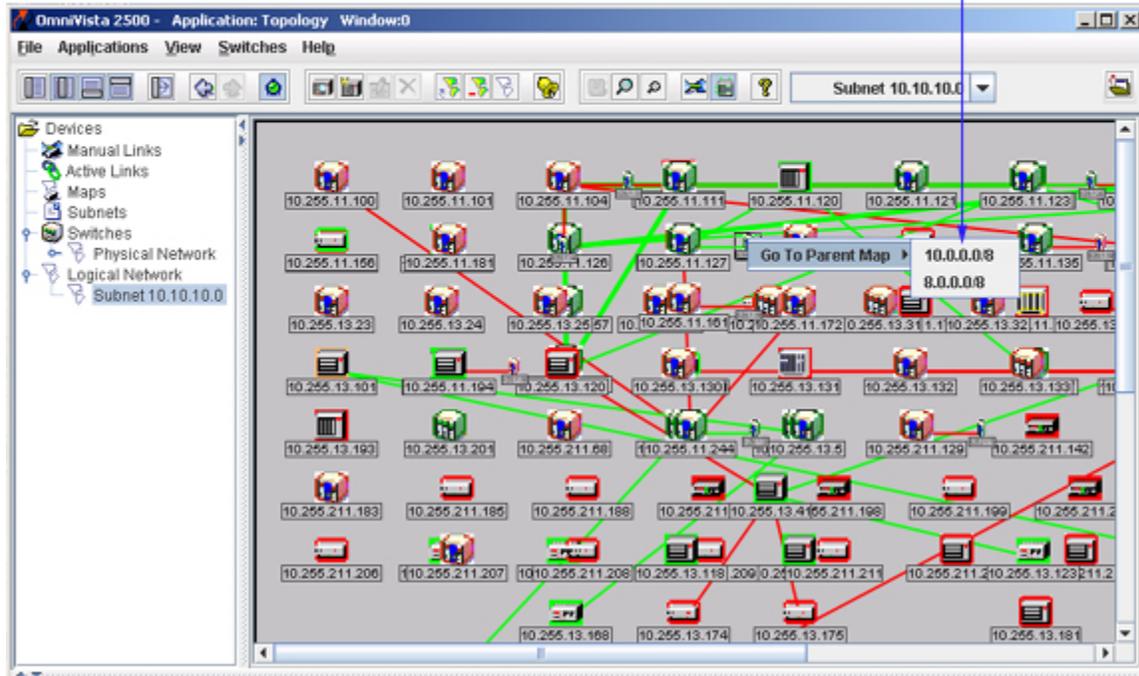


Viewing External Regions

Whenever links to external devices are displayed (by toggling the View External Links Icon ) , you can click right on any external device and select **Go To Parent Map**. When selected, a list of regional maps that contain the selected device display, as shown below. Select the desired regional map to view it. (Note that an "external device" is not part of the regional map displayed. However, it is connected to a device that is part of the regional map displayed.)

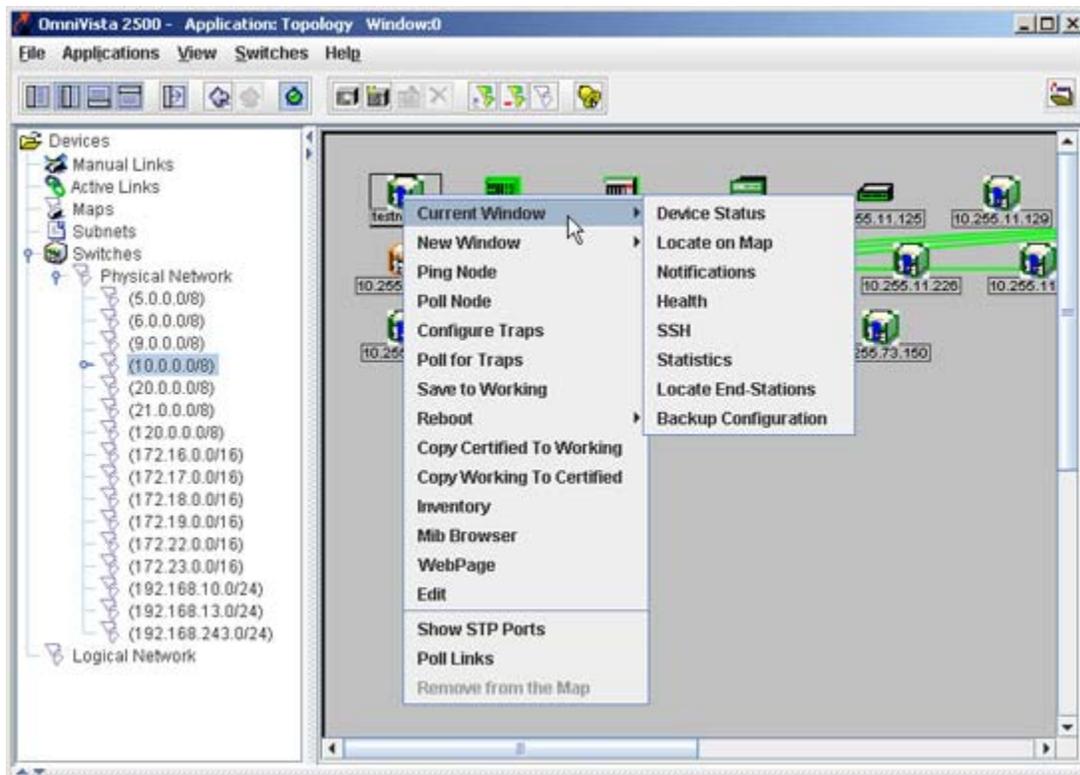
Viewing External Devices

Click right on any external device and select **Go To Parent Map**. All regional maps that contain the selected device display for your selection. Select the desired region to view it.



Pop-Up Menus in Maps

Whenever a regional map is displayed, you can right-click right any device to display a pop-up menu. Different versions of the pop-up menu are displayed for AOS devices, XOS devices, and third-party devices. The first two items on the pop-up menu, **Current Window** and **New Window**, each expand to multiple menu items. **Current Window** and **New Window** enable you to open their respective menu items in the current OmniVista window or in a new, additional OmniVista window. Each menu item displayed on the pop-up menu in a regional map is explained below.



Current Window or New Window > Device Status

Causes OmniVista to select the switch in the Tree and establish a connection to the switch, exactly as if you had manually selected the switch in the Tree. If the switch's icon is not visible in the Tree, OmniVista will expand the Tree and scroll until the switch icon is visible. When a connection is established, device-specific configuration and statistics information displays. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Locate on Map

Loads and displays a regional map in the Physical Network that contains the selected device. The device is automatically selected and centered in the map display. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Notifications

Loads the Notifications application for the selected switch. The Notifications application enables you to view traps for the switch. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Health

Loads the Health application for the selected switch. The Health application displays information on the health of the selected switch. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Telnet or SSH

Either **Telnet** or **SSH** (Secure Shell) displays by default on the pop-up menu, as user-configured for the individual switch. You can configure the default selection for a switch through any of the methods described below. You can also define the switch's Telnet user name and password to OmniVista by means of these methods. When the Telnet user name and password are known, OmniVista will auto login for your

convenience when Telnet or SSH sessions are established. Configure the defaults for a switch using any one of the following methods:

Discover the switch with an SNMP setup that has its **Shell Preference** field set to **Telnet** or **SSH**, as desired. Enter the Telnet user name and password in the respective fields on the SNMP Setups window. (For more information, refer to the help for the Discovery application.)

Edit the switch after discovery and activate the **Prefer SSH** checkbox on the General Tab of the Edit Discovery Manager Entry window. This will specify that SSH is the default for the switch. Enter the Telnet user name and password in the respective fields.

Activate the **Prefer SSH** checkbox on the New Discovery Manager Entry window when you add a switch manually. This will specify that SSH is the default for the switch. Enter the Telnet user name and password in the respective fields.

The **Telnet** or **SSH** menu item opens the Telnet application and establishes a Telnet or SSH connection, respectively, with the selected switch. If the switch's Telnet user name and password are known to OmniVista, auto login will occur. Otherwise you will need to manually enter the switch's Telnet user name and password. Each time the **Telnet** or **SSH** menu item is selected, a new Telnet or SSH session is established. Individual Telnet and SSH sessions are identified by tabs that display the switch IP address. Telnet or SSH sessions can be established in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Statistics

Loads the Statistics application with the Add Item window open and the relevant switch selected automatically. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Locate End-Stations

Loads 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 tab. This function can be performed in the current OmniVista window or in a new OmniVista window.

Current Window or New Window > Backup Configuration

Loads the Backup Configuration utility in the Health application for the selected switch. The Backup Configuration utility in the Resource Manager application loads and saves firmware files for the selected switch. This function can be performed in the current OmniVista window or in a new OmniVista window.

Ping Node

Causes an immediate ping to the selected switch. The result of the ping -- an "equipment is alive" message or an "equipment does not respond" message -- is reported in the Status Panel.

Poll Node

Causes an immediate poll of the selected switch. The success or failure of the poll is reported in the Status Panel.

Configure Traps

Opens the Configure Traps Wizard for the selected switch. The Configure Traps Wizard enables you to configure traps for the switches.

Poll for Traps

Causes an immediate poll for traps of the selected switch.

Save to Working (AOS Devices)

Saves the primary CMM's current running configuration to the working directory of the switch. Executing this command is the same as executing the Save To Working command for an individual device. Click [here](#) for an overview of the directory structure of AOS devices.

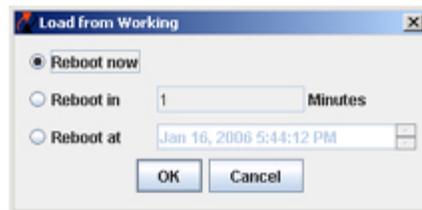
Note: When you apply the **Save to Working** option on a device, you must allow 120 seconds of time to elapse, before you perform the same again.

Reboot > From Working (AOS Devices)

Reboots the primary CMM from the working directory. Executing this command is the same as executing the Load From Working command for an individual device. Note that any unsaved configuration changes will be lost: you can save configuration changes with the **Save to Working** command before executing **Reboot**. Click [here](#) for an overview of the directory structure of AOS devices.

When you select **Reboot > From Working**, the Load from Working window displays. The Load from Working window is shown below. This window enables you to specify whether you wish to reboot immediately (**Reboot now**), or reboot within 1 - 1000 minutes (**Reboot in x Minutes**), or reboot at a specified date and time (**Reboot at date time**). Specify the desired reboot time and then click the **OK** button.

The Load from Working window enables you to schedule the reboot.

**Reboot > From Certified** (AOS Devices)

Reboots the primary CMM from the certified directory. Executing this command is the same as executing the Load From Certified command for an individual device. Note that any unsaved configuration changes will be lost: you can save configuration changes with the **Save to Working** command before executing **Reboot**. Click [here](#) for an overview of the directory structure of AOS devices.

When you select **Reboot > From Certified**, the Load from Certified window displays. The Load from Certified window is shown below. This window enables you to specify whether you wish to reload an entire switch (**Reload Entire Switch**), reboot immediately (**Reboot now**), or reboot within 1 - 1000 minutes (**Reboot in x Minutes**), or reboot at a specified date and time (**Reboot at date time**). Specify the desired reboot time and then click the **OK** button.

The Load from Certified window enables you to schedule the reboot.



Note: When you reboot the primary CMM from the certified directory, the switch will automatically failover to the secondary CMM (in other words, the two CMMs will trade primary and secondary roles). When you reboot the primary CMM from the working directory, no failover occurs.

Copy Certified to Working (AOS Devices)

Copies the contents of the certified directory in the primary CMM to the working directory in the primary CMM. Executing this command is the same as executing the Copy Certified to Working command for an individual device.

Copy Working to Certified (AOS Devices)

Copies the contents of the working directory in the primary CMM to the certified directory in the primary CMM, in a manner similar to the **Copy Certified to Working** command described above.

Note: The **Copy Working to Certified** command also automatically synchronizes the switch's CMMs after the copy operation is completed.

Inventory

Loads the Inventory application for the selected switches. The Inventory application enables you to create reports. The reports can include system information, detailed module information, chassis information, and health information.

MIB Browser

Loads the OmniVista MIB Browser for the selected switch.

WebPage or SwitchManager or TrackView

This menu item opens the device manager that is appropriate for the selected switch. WebView, the Alcatel device manager, opens for AOS devices. WebView enables you to perform direct device-level AOS configuration from a browser. **TrackView** opens for OmniCore devices. **WebPage** opens for the OmniStack 1024, 6024, 6300-24, and 8008, as well as the OmniMSS. **SwitchManager** opens for all other XOS devices. Each device manager enables you to perform device-level configuration of the selected device.

Note: SwitchManager and TrackView will open only if the respective program is installed on the client.

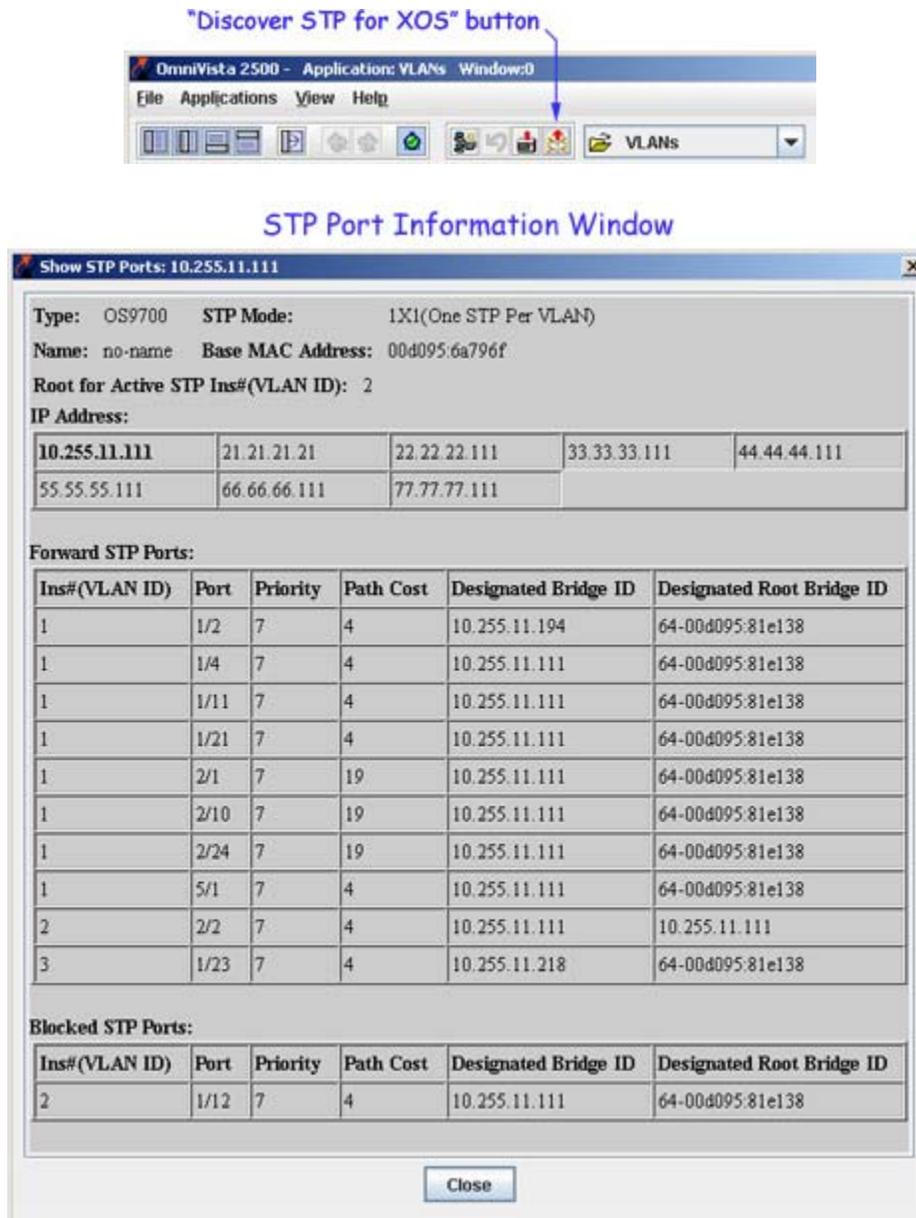
Edit

Opens the Edit Discovery Manager Entry window, which enables you to edit devices. When you edit a device, it is important to understand that you are editing OmniVista's knowledge of the device, not the device itself. Click here for the reasons you might want to edit a device.

Show STP Ports (AOS and XOS Devices Only)

When you right-click on a switch, a window opens showing STP port information collected for the selected switch. In addition, on AOS devices, STP information for LAG ports is displayed. If the switch is configured for MSTI, only non-MSTI information is displayed. You must have "Write" permission to perform this function.

Note: To display STP information for XOS switches, you must first initiate an STP poll by clicking on the **Discover STP for XOS** button in the VLANs Application. A red line appears on the button when you start polling. When polling is complete, the line disappears. Clicking again on the button before polling is complete, will stop the polling.



Type - The switch model number (e.g., OS6624).

Name - The user-configured name for the switch.

STP Mode - The STP operating mode for the switch (e.g., Flat, 1x1).

Base MAC Address - The System MAC address of the switch.

Root for Active STP Instance -The VLAN(s) for which this switch is the Root Bridge.

IP Address - The IP address(es) available on the switch.

Port Information -

Instance # (VLAN ID) - The VLAN ID(s) associated with the VLAN Spanning Tree instance.

Port - The slot number for the module and the physical port number or a logical port. If the slot number is 0, the port number refers to a link aggregate logical port number.

Priority - The Spanning Tree priority for the port (0–15). The lower the number, the higher the priority.

Path Cost - The contribution of this port to the path cost towards the Spanning Tree root bridge that includes this port. Path cost is a measure of the distance of the listed port from the root bridge in the number of hops.

Designated Bridge ID - The IP or MAC address of the designated bridge for this port's segment.

Designated Root Bridge ID - The IP or MAC address of the Designated Root Bridge for the port. If the port is an Edge Port, the field will display ffff-fffff.fffff.

Poll Links

Causes an immediate poll of all links associated with the selected device to gather current information on link status. The success or failure of the link poll is reported in the Status Panel. Note that the Poll Links icon enables you to poll links on multiple selected devices, or on all devices present in the map. Click here for further information on the Poll Links icon.

Remove from the Map (active in the Logical Network only)

Deletes the selected device from the region in the Logical Network that is currently displayed. However, it does not remove the device from the Physical Network nor from the list of All Discovered Devices. (To remove a device from the Physical Network, and from the list of All Discovered Devices, select the device in the List of All Discovered Devices, click right, and select **Remove** from the pop-up menu that displays.)

Device Label Options

You can select the information used for device labels in the Tree and in regional maps. You can select IP address only, device name only, or both. To specify the device labels you want, go to the View menu and select **Devices By**, as shown below. Then select the desired option from the submenu displayed. Your selection is effective immediately.



Note: The device labels you specify are used throughout the Topology application AND other applications such as Notifications.

Managing Regional Maps

The Topology application enables you to graphically view a map of the overall Physical Network or the overall Logical Network. You can also view a map of any individual region (i.e., subnet) in the Physical Network or any individual region in the Logical Network. To view a regional map, select the region that you want to view in the Tree. You can select the Physical Network, the Logical Network, any subnet in the Physical Network, or any region in the Logical Network.

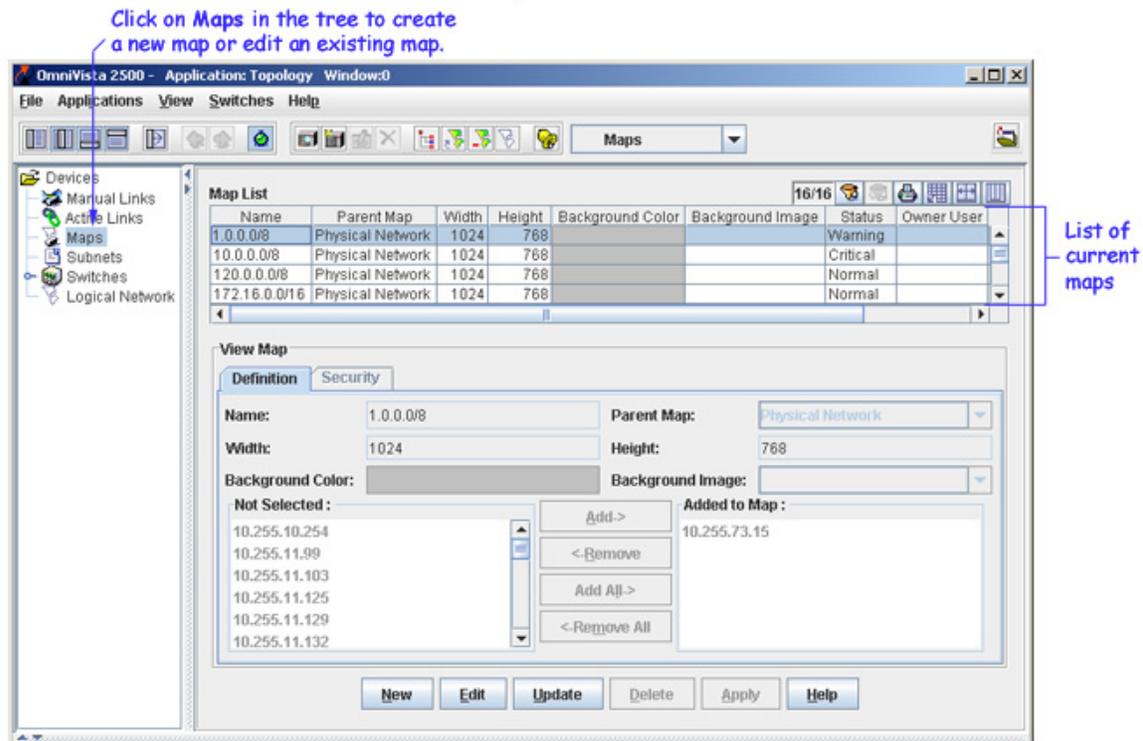
The Physical Network, as its name implies, is an image of the physical subnets and devices in the network. When OmniVista discovers the network, it arranges the discovered devices into default subnets. You can override OmniVista's default subnet creation by creating manual, that is, user-defined, subnets. (Click here for more information on default and manual subnet creation.) However, all subnets in the Physical Network, both default subnets and manual subnets, are created according to the device IP address. You cannot "pick and choose" the individual devices to be included in a subnet.

In contrast, within the Logical Network you can create "logical regions" and select the individual devices to be included in the region, regardless of the device IP address. You can create logical regions where devices are grouped and displayed in any way that is meaningful for your individual network, in any configuration desired.

The Maps Window

The Maps window, shown below, enables you to create regional maps in the Logical Network from scratch, create regional maps in the Logical Network from existing subnets in the Physical Network, edit existing regional maps in both the Logical Network and the Physical network, and delete regional maps from the Logical Network or the Physical Network. When you create or edit a regional map, you can define the background color you want used when the map is displayed and the width and height of the viewing window. You can also specify a background image for the map, if desired.

The Maps Window



Creating a New Map From Scratch

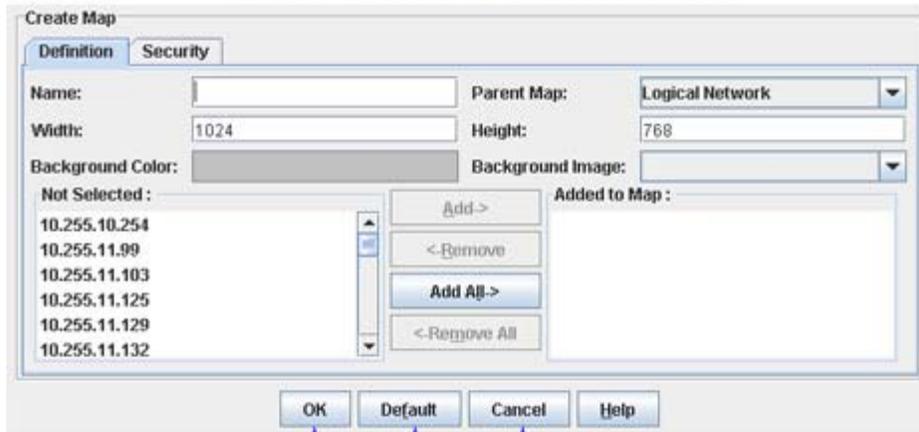
To create a new regional map in the Logical Network from scratch, click the **New** button, which is visible in the screen above. The **Create Map** panel activates, as shown below. The Definition tab is used to create the map. The Security tab is used to configure viewing permissions for the map.

Note: You can create a map without defining viewing permissions. The viewing permissions will be set to the default of "Owner", and only the person who created the map will be able to view it..

Definition Tab

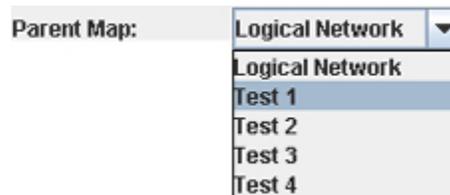
Follow the steps below to create a new map.

Create Map Panel - Definition Tab



Create a new regional map in session memory with the values displayed.
 Reset all fields to default values.
 Deactivate the Create Map Panel.

1. Enter a name for the new map in the **Name** field.
2. Define the region that is the "parent" of the new regional map by selecting a region from the **Parent Map** drop-down menu, shown below. All existing regional maps in the Logical Network are listed for your selection. All regional maps in the Logical Network must have a parent region defined; the default parent region is **Logical Network**.



3. Define the size of the viewport that will display the new regional map by entering the desired width (in pixels) in the **Width** field and the desired height (in pixels) in the **Height** field.
4. Define a background color for display of the new regional map. To do this, click left anywhere in the color displayed by the **Background Color** field. The Color Chooser displays. The Color Chooser enables you to define a background color by selecting a color swatch from a group of pre-defined swatches, or by using the HSB (hue saturation brightness) color model, or by using the RGB (red green blue) color model.
5. If you want the new regional map to display against a background image, select an image from the **Background Image** drop-down menu. All background images must be imported into OmniVista. Importing background images is explained below in Importing Background Images for Regions.
6. Define the devices that are part of the new regional map by selecting switches in the **Not Selected** area and moving them into the **Added to Map** area using the **Add >**, **< Remove**, **Add All >**, and **< Remove All** buttons. Continue moving devices until the **Added to Map** area contains all devices that you want placed in the new regional map.

7. Click the **OK** button. The new regional map is written to session memory and is added to the Maps List. However, the new regional map is not yet written to the server: it is an "unsaved" change.

8. Click the **Apply** button to save the new regional map to the server.

Note: To set viewing permissions, complete the fields in the **Security** tab (described below) before clicking the **Apply** button.

Security Tab

After creating a map, you can set permissions on who can view the map. You can create a map without defining viewing permissions. The viewing permissions will be set to the default of "Owner", and only the person who created the map will be able to view it.. To set viewing permissions, complete the fields in the **Security** tab as described below.

1. Click on the **Security** tab.



2. The **Owner User** Field is pre-filled (e.g., admin).

3. Click on the drop-down menu in the **View Permission** field to set viewing permissions.

Owner User - Map is visible to only the owner and Network Administrator.

Group Users - Map is visible to the owner, all users of the groups he belongs to, and Network Administrator.

All Users - Map is visible to all.

4. Click the **OK** button. The new regional map is written to session memory and is added to the Maps List. However, the new regional map is not yet written to the server: it is an "unsaved" change.

5. Click the **Apply** button to save the new regional map to the server.

Creating New Maps from Subnets

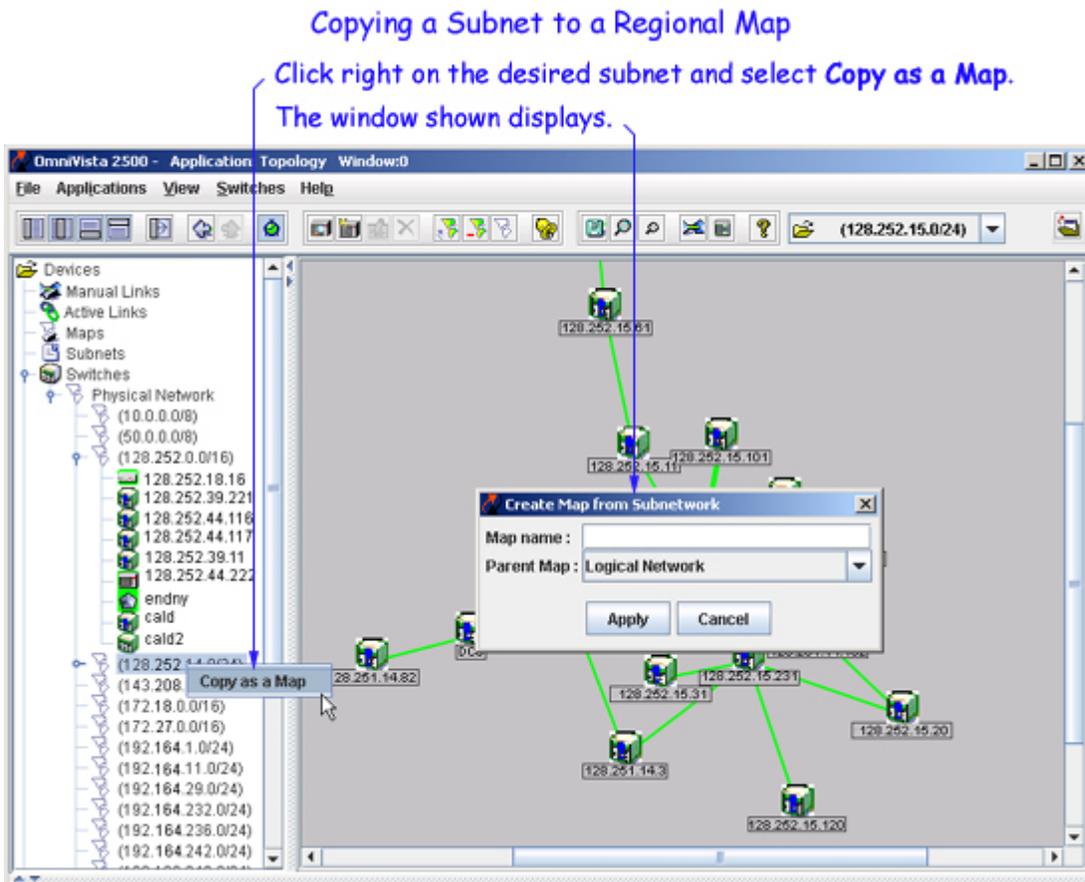
You can create a new regional map in the Logical Network from one or more subnets in the Physical Network. When you do this, all switches in the physical subnet(s) are automatically placed into the new logical map. There are two ways to create regional maps in the Logical network from physical subnets:

Right-click on any physical subnet listed in the Tree and "copy" it to a new regional map in the Logical Network.

Use the "Create Map from Subnets" Wizard to create a new regional map in the Logical Network from one or more subnets in the Physical Network. The wizard also gives you the option of filtering the subnets in the map by VLAN (i.e., a device will be included if the selected VLAN exists within it).

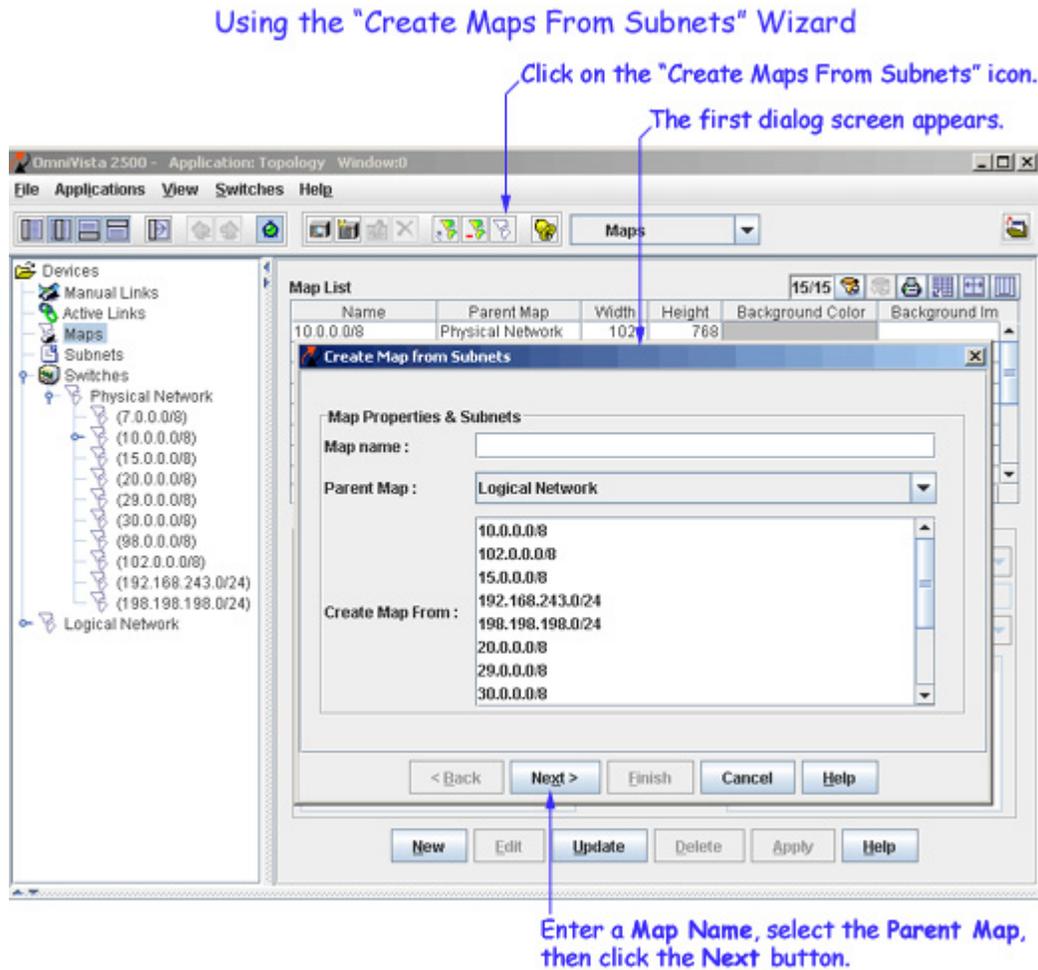
Copying a Subnet to a Regional Map

To copy an individual physical subnet to a new regional map in the Logical Network, right-click on the desired subnet and select the **Copy as a Map** menu item, as shown below. The window shown displays. Enter a name for the new map in the **Map Name** field and select the parent map from the **Parent Map** drop-down menu, then click the **Apply** button. The new regional map is written directly to the server and displays in the Maps List.



Using the "Create Maps from Subnets" Wizard

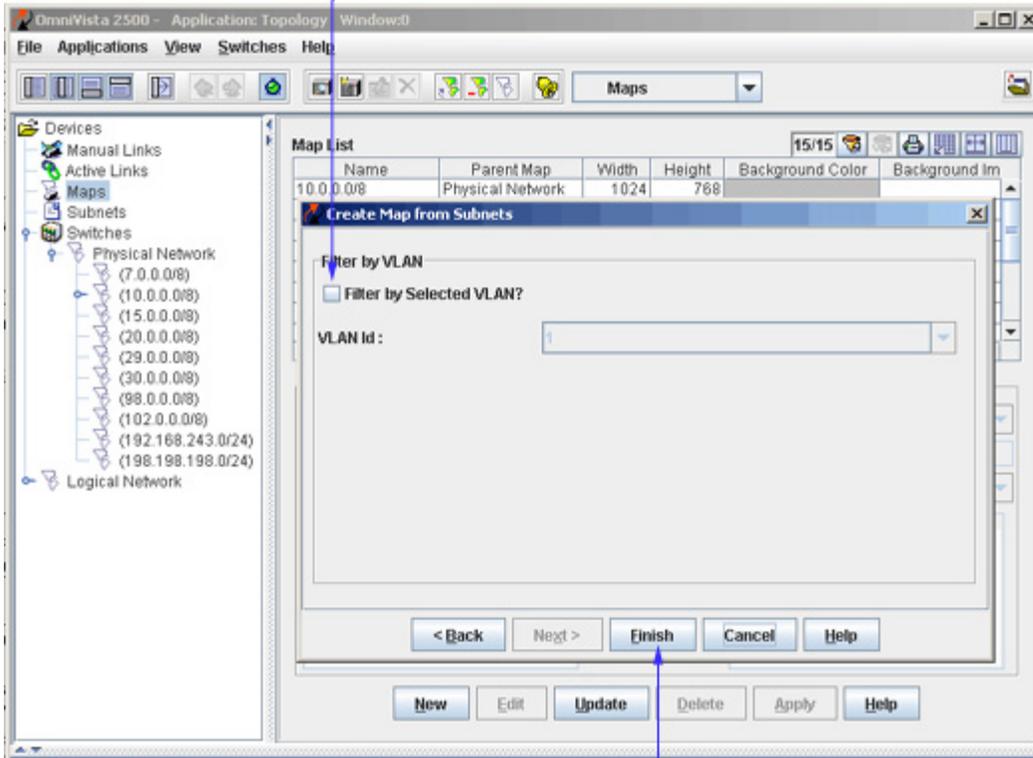
Click on the "Create Maps from Subnets" icon to bring up the first page of the wizard. Enter a name for the new regional map in the **Map Name** field and select the parent map from the **Parent Map** drop-down menu. In the **Create Map From** area, select the physical subnet(s) that you want included in the new map. You can select multiple contiguous subnets by **Shift**-clicking and non-contiguous subnets by **Ctrl**-clicking. Click the **Next** button.



On the second screen of the wizard, you have the option of creating the map to include all devices in the map or filtering the map by VLAN. To complete the map without filtering, just click on the **Finish** button. To filter the devices in the map by VLAN, click on the **Filter by Selected VLAN** checkbox to activate the **VLAN ID** drop-down menu, then select the VLAN.

Using the "Create Maps From Subnets" Wizard (Cont.)

To filter the map by VLAN, select the checkbox then select the VLAN from the drop-down menu.



Click the Finish button.

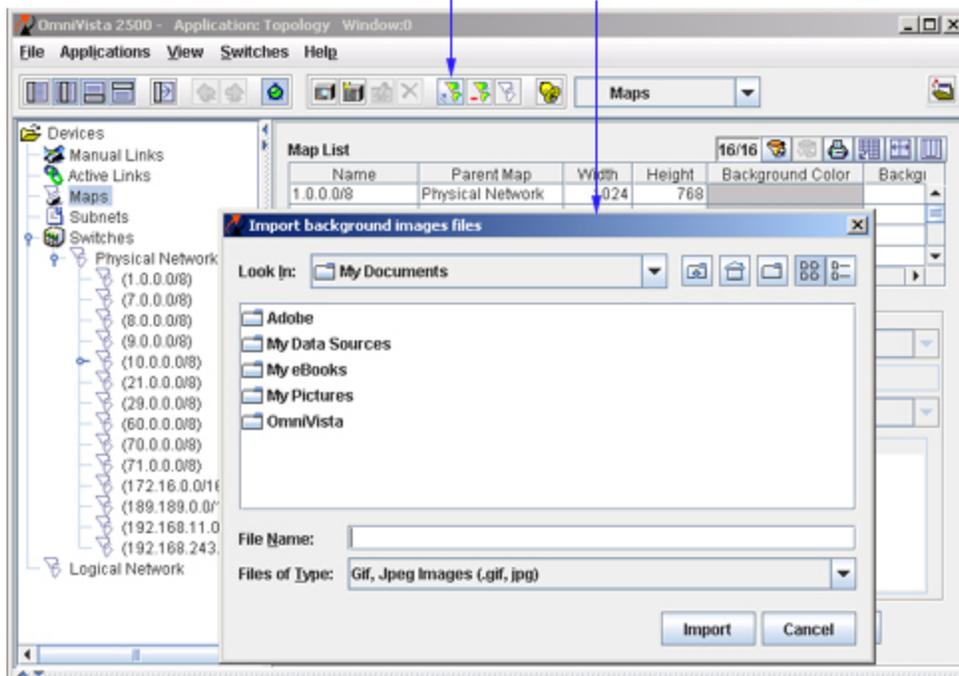
When you click the **Finish** button, the new regional map is written directly to the server and displays in the Map List. The "Filter by VLAN" Map behaves the same as any other map in OmniVista.

Importing Background Images

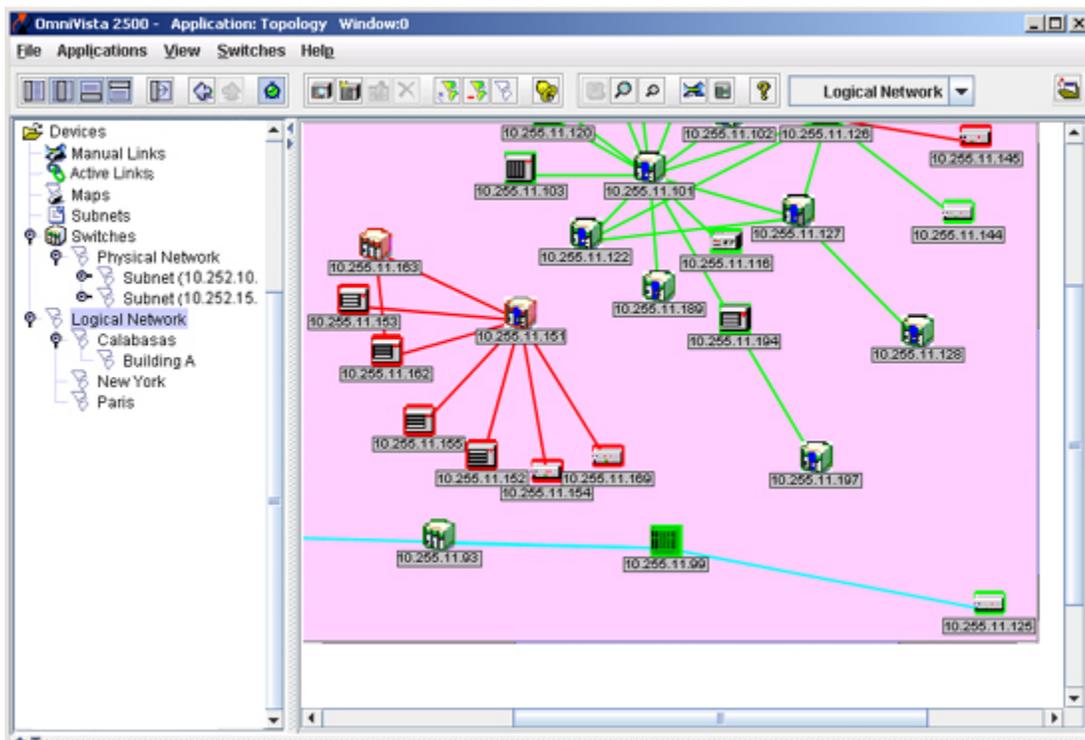
When you create a new regional map from scratch (or when you edit an existing regional map) you can specify a background image to be used when the map displays. Background images can be Gif (.gif) files or Jpeg (.jpg) files. You must import any background images you want to use. To import image files, click the **Import Backgrounds** icon, as shown below, or select **Import Backgrounds** from the **File** menu. Locate the files you want to import, then click the **Import** button. All background image files that you import will automatically display in the **Background Image** drop-down menu. Note that you can delete imported background images by clicking the **Remove Backgrounds** icon  or by selecting **Remove Backgrounds** from the **File** menu.

Importing Background Image Files

Click the "Import Backgrounds" icon
The window shown displays



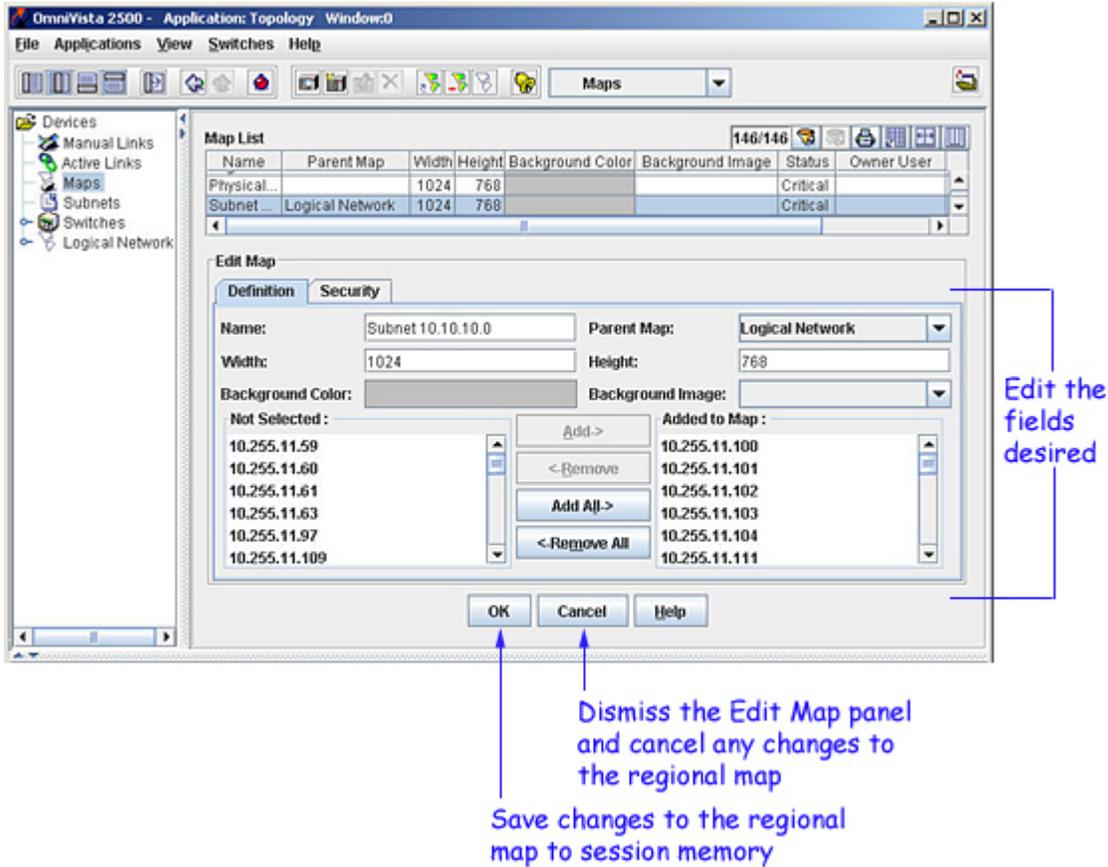
Example of a Map with a Background Image



Editing a Map

To edit an existing regional map, select the map in the Maps List and click the **Edit** button. The **Edit Map** panel activates, as shown below. The Definition tab is used to modify the map definitions (e.g., Name, Colors, Switches). The Security tab is used to modify viewing permissions for the map.

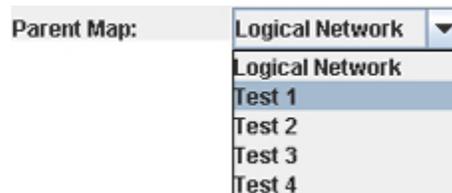
Editing a Regional Map



Definition Tab

Modify the fields in the definition tab as described below.

1. Modify the name for the map in the **Name** field.
2. Modify the region that is the "parent" of the new regional map by selecting a region from the **Parent Map** drop-down menu, shown below. All existing regional maps in the Logical Network are listed for your selection. All regional maps in the Logical Network must have a parent region defined; the default parent region is **Logical Network**.



3. Modify the size of the viewport that will display the new regional map by entering the desired width (in pixels) in the **Width** field and the desired height (in pixels) in the **Height** field.
4. Modify a background color for display of the new regional map. To do this, click left anywhere in the color displayed by the **Background Color** field. The Color Chooser displays. The Color Chooser enables you to define a background color by selecting a color swatch from a group of pre-defined swatches, or by using the HSB (hue saturation brightness) color model, or by using the RGB (red green blue) color model.
5. To modify the background image, select an image from the **Background Image** drop-down menu. All background images must be imported into OmniVista. Importing background images is explained below in Importing Background Images for Regions.
6. Add or delete switches from the map by selecting switches in the **Not Selected** area and moving them into the **Added to Map** area using the **Add**, **<Remove**, **Add All**, and **<Remove All** buttons. Continue moving devices until the **Added to Map** area contains all devices that you want placed in the new regional map.
7. Click the **OK** button. The changes are written to session memory.
8. Click the **Apply** button to save the changes to the server.

Note: To modify viewing permissions, modify the fields in the **Security** tab before clicking the **Apply** button.

Security Tab

If the map is a user-defined logical map, the **Security** tab will be enabled for the user to edit permissions for viewing the map. By default, the "Owner" of the map (the user who created the map) is allowed to view the map. However, you can expand the viewing permissions using the **Security** tab.



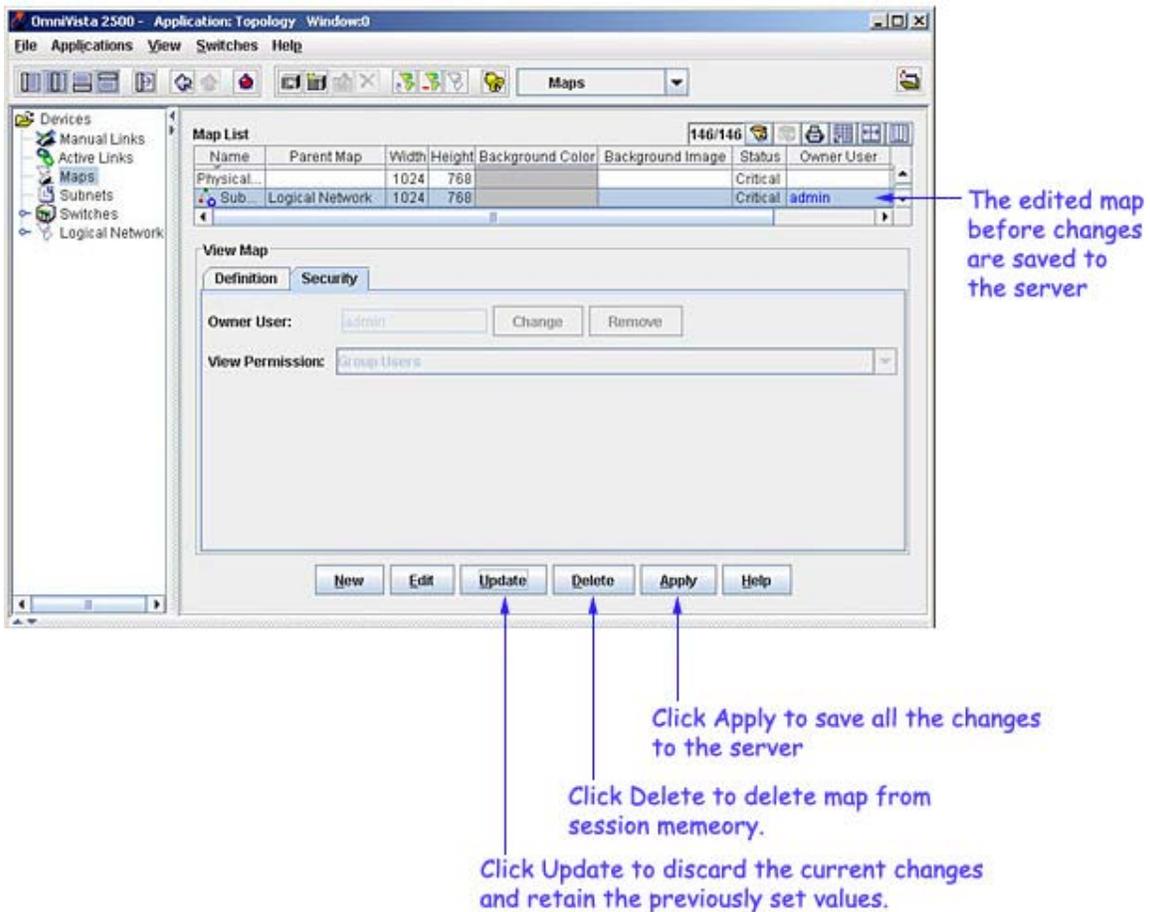
Changing the Ownership of the Map

1. To change the owner of the map, click the **Change** button next to the **Owner User** field, select the new Owner User from the **Select User** window, then click **OK**. (Only users who can modify the map are listed.)



Note: To remove the ownership for the map, click **Remove**, then click **Yes** at the confirmation prompt.

2. Click the **OK** button at the bottom of the panel to save the changes. The Maps window will be displayed with the edited maps as shown below.



3. Click the **Apply** button to save the changes to the server.

The following points need to be remembered, when changing the ownership of the map or setting permissions for viewing the map.

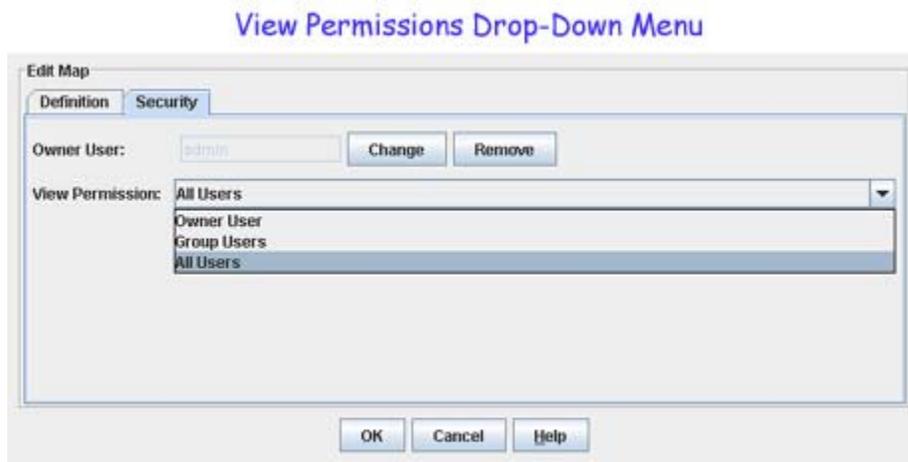
Deleting the ownership of a map or changing its view permissions will affect the ownership and permission of all maps in the sub-tree below this map.

If deleting the **Owner User** of a map or setting **View Permissions** causes the child map to have higher visibility than its parent, a warning message will be displayed prompting the user to change the current settings.

Changing the View Permission for the Map

The **View Permissions** field determines which users can see the active map. Modify the permissions as described below.

1. Select the permission level from the **View Permission** drop-down menu.



None - The map is visible to only the owner and Network Admin.

Group Users - The map is visible to the owner, all users of the group he belongs to and to the Network Admin.

All Users - The map is visible to all.

2. Click the **OK** button.

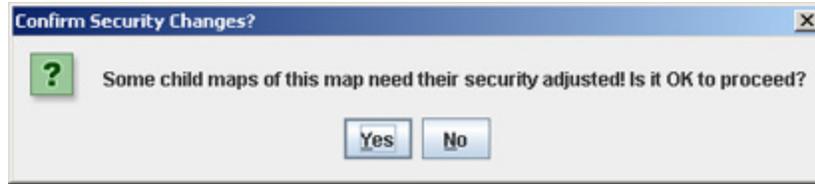
3. Click the **Apply** button to save the changes to the server.

The following points need to be remembered, when changing the ownership of the map or setting permissions for viewing the map.

Deleting the ownership of a map or changing its view permissions will affect the ownership and permission of all maps in the sub-tree below this map.

If deleting the **Owner User** of a map or setting **View Permissions** causes the child map to have higher visibility than its parent, a warning message will be displayed prompting the user to change the current settings.

If a parent map's owner is changed, the following prompt may appear:

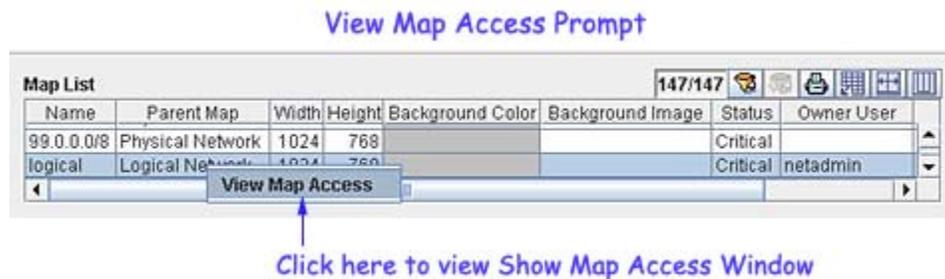


Select the **Yes** button to change the user and fix all child maps, as necessary.

If a map is visible to a user who is a Writer and not its owner, the user can modify the map but not its ownership or parent, i.e. he can modify its size, background color, image, or devices in the map.

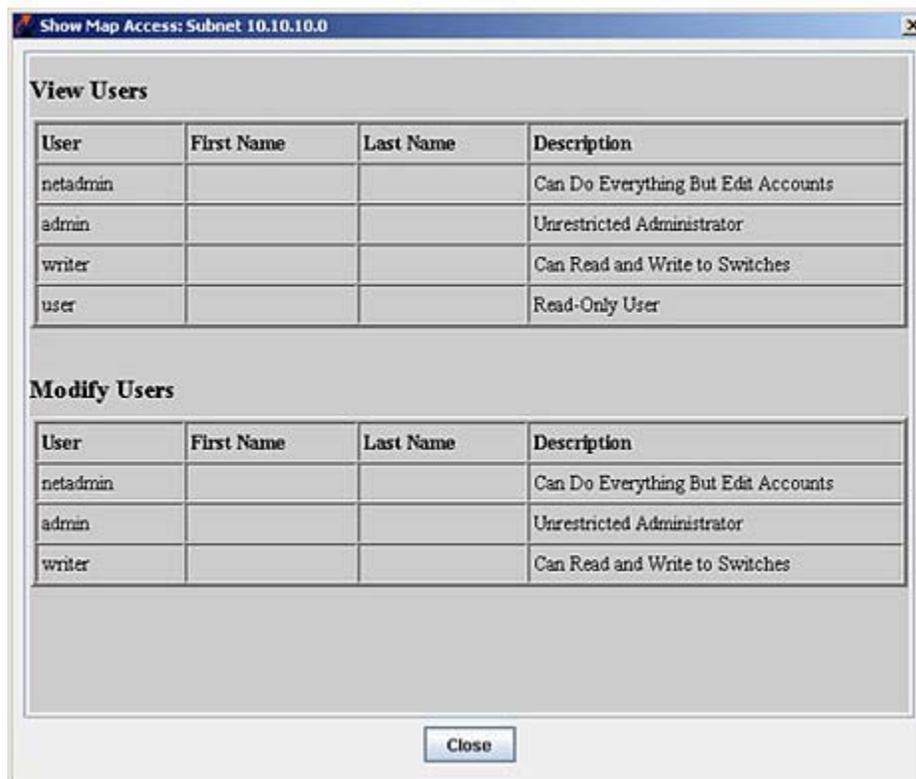
Viewing Map Ownership/Viewing Permissions

To view map properties for a user-created logical map, right-click on the map in the Map List. The **View Map Access** prompt will be displayed.



Click on the prompt to view the **Show Map Access** Window.

Show Map Access Window



Deleting a Map

To delete an existing regional map, follow the steps below:

1. Select the map in the Maps List.
2. Click the **Delete** button. The map is deleted from session memory..
3. Click the **Apply** button to delete the map from the server.

Using the Color Chooser

The Color Chooser has three tabs that enable you to define a color in any of three different ways:

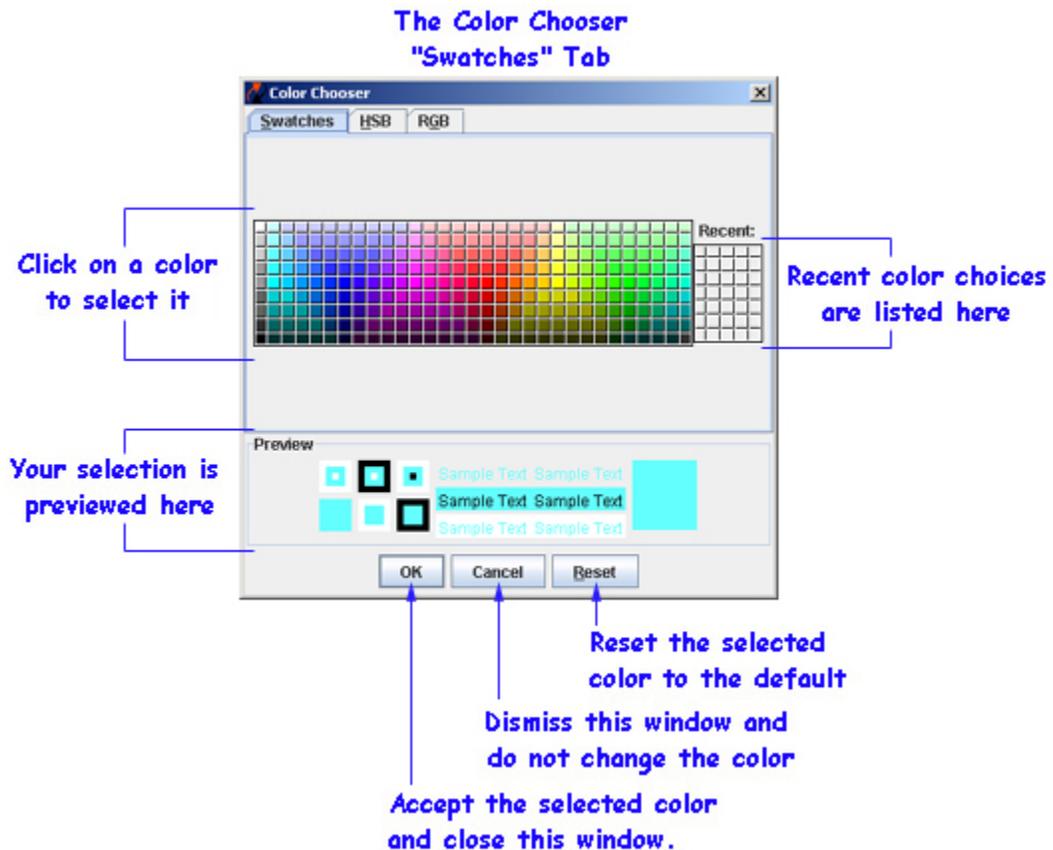
The Swatches Tab enables you to select a color swatch from a group of predefined swatches.

The HSB Tab enables you to define a color using the HSB (hue saturation brightness) color model.

The RGB Tab enables you to define a color using the RGB (red green blue) color model.

The Swatches Tab

The Swatches Tab, shown below, enables you to select a background color from predefined color swatches.



The HSB Tab

The HSB Tab, shown below, enables you to define a color using the HSB color model. HSB uses three axes to define a color: hue, saturation, and brightness. *Hue* defines the color itself -- for example, red instead of blue or yellow. *Saturation* defines the degree to which the hue differs from neutral gray. Saturation values can range from 0, which means no color saturation, to 100, which means the fullest

saturation of a given hue at a given brightness. *Brightness* defines the level of illumination. Brightness values can range from 0, which appears black (as there is no light) to 100 for full illumination, which appears white (as all the color is washed out).

The Color Chooser
HSB (Hue, Saturation, Brightness) Tab

To select a color:
Click on a color to select it. The color's HSB and RGB values display and the color is previewed. You can also enter HSB values directly into the "H", "S", and "B" fields.

To adjust the colors displayed for selection:
Click "H" to adjust the hue, "S" to adjust the saturation, "B" to adjust the brightness. Perform the adjustment by sliding this lever. Equivalent RGB values are displayed for reference.

Reset the selected color to the default

Dismiss this window and do not change the color

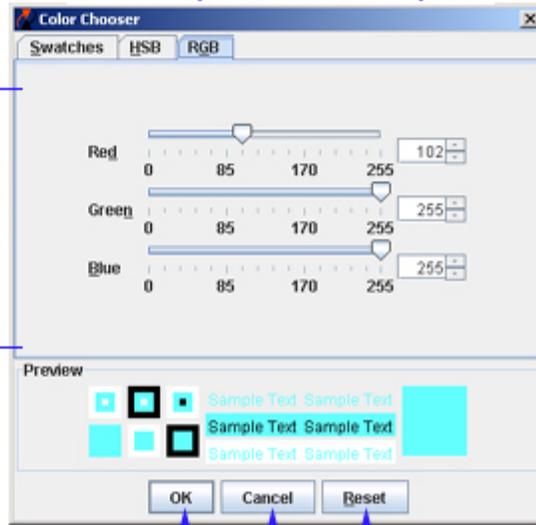
Accept the selected color and close this window.

The RGB Tab

The RGB Tab, shown below, enables you to define a color using the RGB color model. RGB defines a color by specifying the individual amounts of red, green, and blue to be added to the color.

The Color Chooser The RGB (Red, Green, Blue) Tab

Slide the red, green, and blue controls to define the desired color. Colors are previewed as you slide the controls.



Reset the selected color to the default

Dismiss this window and do not change the color

Accept the selected color and close this window.

Viewing Active Links

OmniVista now includes the ability to view a tabular listing of the active links in the network. The list includes all links that were learned during the discovery process and all links that were created manually or imported into OmniVista. To view a list of active links, merely click **Active Links** in the tree, as shown below. Each field in the list of active links is described below.

Note: When you click the **Update** button the active links in the local database and not necessarily the current active links will be displayed. Perform a polling operation to update the local database.

Viewing Active Links

Click Active Links in the Tree

The list of Active Links displays

| Origin | IP Address 1 | Slot/Port 1 | LAG 1 | IP Address 2 | Slot/Port 2 | LAG 2 | Media Type | Status | VLAN Id | Discovered |
|--------|---------------|-------------|-------|---------------|-------------|-------|----------------|--------|---------|------------|
| AMAP | 10.255.11.100 | 2/17 | | 10.255.11.244 | 1/37 | | ethernetCsmacd | Up | 1 | true |
| AMAP | 10.255.11.100 | 2/5 | | 10.255.11.224 | 1/10 | | ethernetCsmacd | Up | 1 | true |
| XMAP | 10.255.11.102 | 8/23 | | 10.255.11.119 | 2/11 | | eth | Up | 1 | true |
| AMAP | 10.255.11.111 | 1/4 | | 10.255.11.102 | 8/14 | | ethernetCsmacd | Up | 50 | true |
| AMAP | 10.255.11.111 | 1/5 | | 10.255.11.102 | 8/6 | | ethernetCsmacd | Up | 1 | true |
| AMAP | 10.255.11.112 | 1/1 | | 10.255.11.102 | 8/3 | | ethernetCsmacd | Up | 1 | true |
| AMAP | 10.255.11.112 | 1/22 | | 10.255.11.111 | 1/11 | | ethernetCsmacd | Up | 70 | true |
| XMAP | 10.255.11.120 | 2/12 | | 10.255.11.207 | 1/24 | | eth | Down | 1 | true |
| AMAP | 10.255.11.121 | 1/3 | | 10.255.11.174 | 2/19 | | ethernetCsmacd | Down | 1 | true |
| AMAP | 10.255.11.126 | 2/12 | | 10.255.11.207 | 1/24 | | eth | Down | 1 | true |
| XMAP | 10.255.11.153 | 2/3 | | 10.255.11.161 | 3/10 | | eth | Up | 1 | true |
| AMAP | 10.255.11.161 | 1/1 | | 10.255.11.157 | 1/11 | | ethernetCsmacd | Up | 1 | true |
| AMAP | 10.255.11.161 | 1/9 | | 10.255.11.125 | 1/18 | | ethernetCsmacd | Down | 1 | true |
| AMAP | 10.255.11.198 | 13/1 | | 10.255.11.111 | 1/24 | | ethernetCsmacd | Up | 40 | true |
| AMAP | 10.255.11.198 | 13/11 | | 10.255.11.112 | 1/23 | | ethernetCsmacd | Up | 70 | true |
| AMAP | 10.255.11.198 | 13/3 | | 10.255.11.102 | 8/10 | | ethernetCsmacd | Up | 1 | true |
| AMAP | 10.255.11.221 | 1/4 | | 10.255.11.224 | 1/14 | | ethernetCsmacd | Up | 1 | true |
| AMAP | 10.255.11.231 | 9/1 | | 10.255.11.224 | 1/15 | | ethernetCsmacd | Up | 1 | true |
| AMAP | 10.255.11.244 | 1/28 | | 10.255.11.172 | 1/34 | | ethernetCsmacd | Up | 1 | true |
| XMAP | 10.255.13.39 | 4/4 | | 10.255.13.38 | 5/4 | | eth | Down | | true |

Click Update to update the screen with the most-recent information

Origin

The origin of the link, which can be **XMAP** (XOS devices), **AMAP** (AOS and OmniStack 61xx and 6300-24) devices, **PNNI** (ATM Private Network-to-Network Interface), **Manual** (manually created), or **Locator** (OmniVista Locator application).

IP Address 1

The IP address of one switch in the link.

Slot/Port1

The slot and port that connect the link on the first switch, specified above.

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 IP address of the second switch in the link.

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.

Media Type

The media type of the link.

Status

The status of the link: **Up**, **Down** or **Unknown**.

VLAN Id

The VLAN Id associated with the AMAP link of the AOS device.

Note: The VLAN Id column for XOS devices will be empty because OmniVista will not display VLAN information for XOS devices.

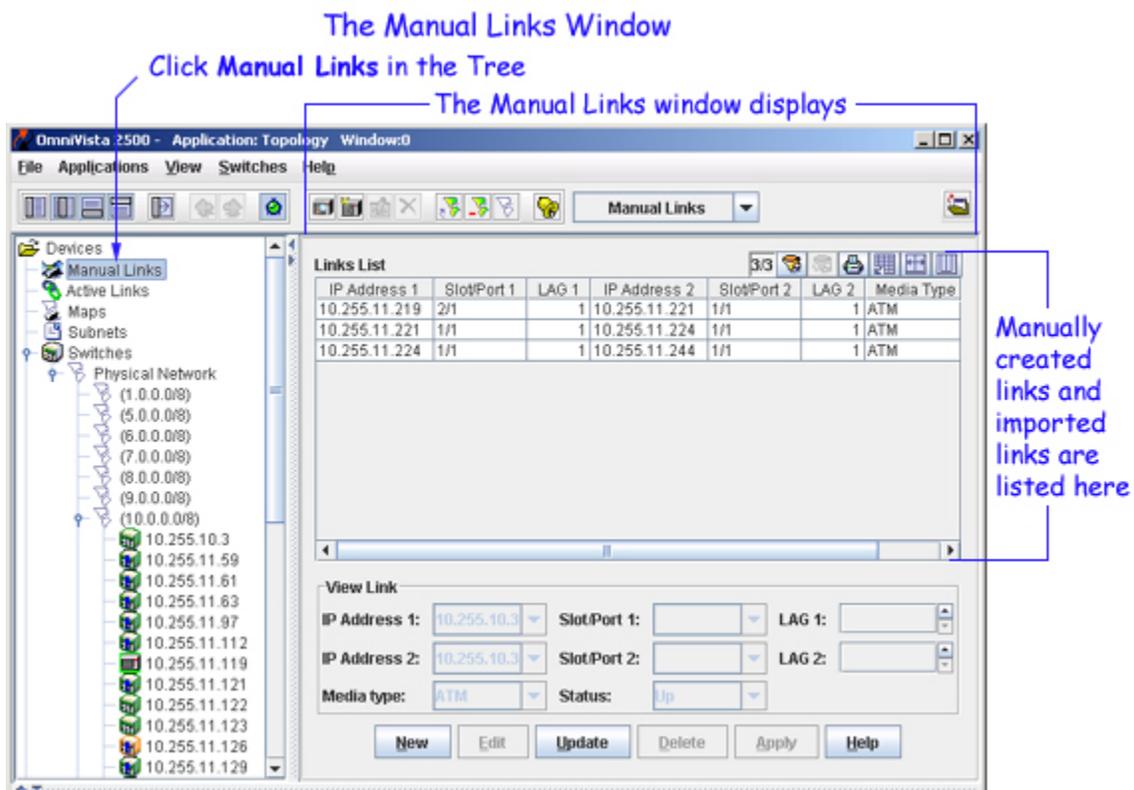
Discovered

This field displays **true** if the link was discovered via the discovery process. This field displays **false** if the link was created manually or imported.

Managing Manual Links

The Manual Links window, shown below, displays a list of the links that were manually created or that were imported into OmniVista. The Manual Links window enables you to create new links manually, to edit existing links that were created manually, and to delete links that were created manually. The Manual Links window also enables you to import links (from a Microsoft Excel file) and to export links (to a Microsoft Excel file).

Note: The Manual Links window does not display links that were learned during the discovery process. Such links are displayed and listed in the Active Links window.



Creating New Links Manually

To create a new link manually, click the **New** button on the Manual Links window. The Create Link panel activates, as shown below. Note that if you select an existing link and then click the **New** button, the Create Link panel will display values from the selected link. This is convenient when you are creating multiple new links.

The Create Link Panel

If you select an existing link and click **New**, the Create Link panel will display values from the selected link.

OK Cancel Help

Create the new link specified.

Dismiss the panel and do not create a link.

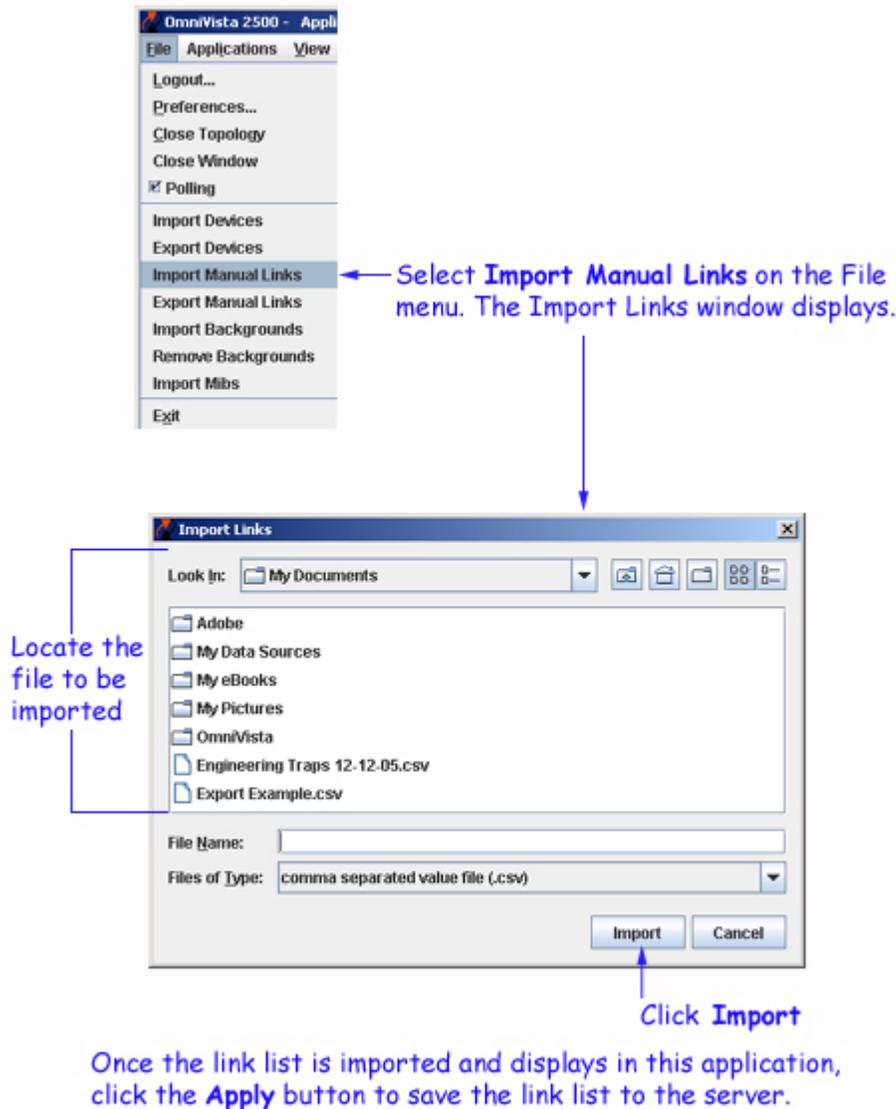
To create a new link, follow the steps below:

1. Set the **IP Address 1** field to the IP address of one switch in the link. All known switches are displayed for your selection.
2. Set the **Slot/Port 1** field to the slot and port that connect the link on the switch specified above. (**Note:** This drop-down bar also displays the port's description found in the MIB table.)
3. If this is a link aggregation link, set the **LAG 1** field to the Link Aggregation reference number assigned by the switch specified above when the link aggregation group was created.
4. Set the **IP Address 2** field to the IP address of the second switch in the link. All known switches are displayed for your selection.
5. Set the **Slot/Port 2** field to the slot and port that connect the link on the second switch. (**Note:** This drop-down bar also displays the port's description found in the MIB table.)
6. If this is a link aggregation link, set the **LAG 2** field to the Link Aggregation reference number assigned to the link aggregation group by the second switch.
7. Set the **Media Type** field to the media type of the link.
8. Set the **Status** field to **Up** or **Down** to define the status of the link. If set to **Up**, the link will display green. If set to **Down**, the link will display red. Note that you can edit the link later if you wish to change its status.
9. Click the **OK** button. The link is created in this application and displays in the Links List.
10. Click the **Apply** button to save the new link to the server.

Note: When you click the **Update** button, the list of all the manually created links in the OmniVista will be refreshed and displayed.

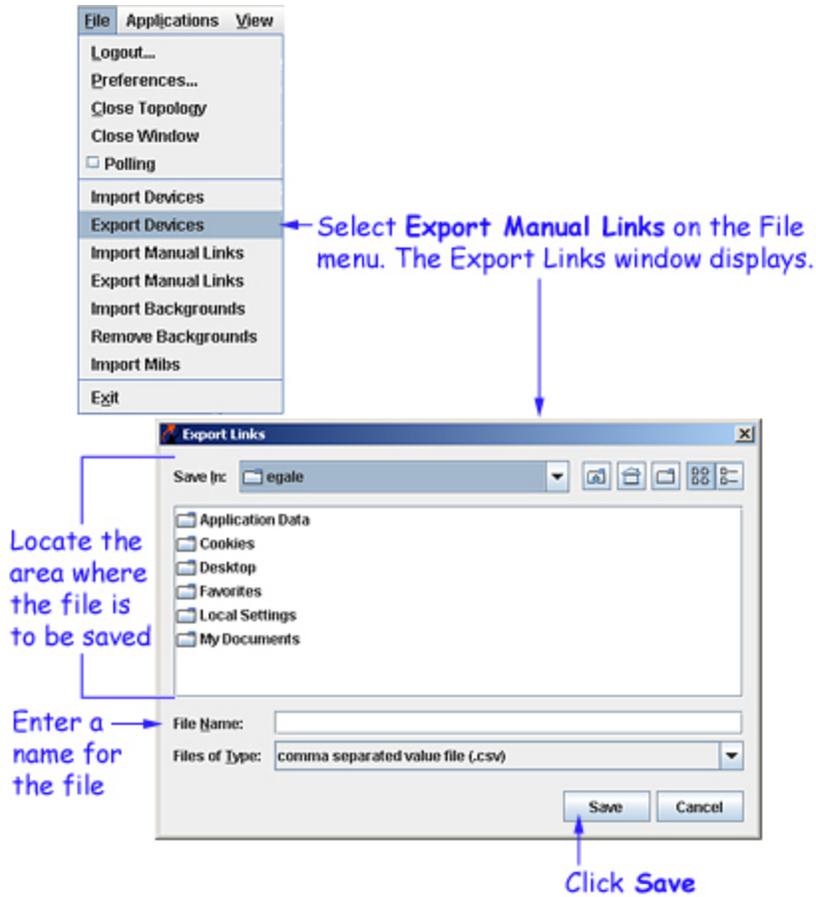
Importing Links

You can import a list of links into this application from a Microsoft Excel file or any other application that produces comma-separated value files (.csv file extension). A comma-separated value file, as the name implies, lists a series of values separated by commas. To import a list of links into this application, select **Import Manual Links** on the File menu. The Import Links window displays, as shown below. Locate the file that you wish to import and then click the **Import** button. When the imported links display in the Manual Links window, click the **Apply** button to save the link list to the server.



Exporting Links

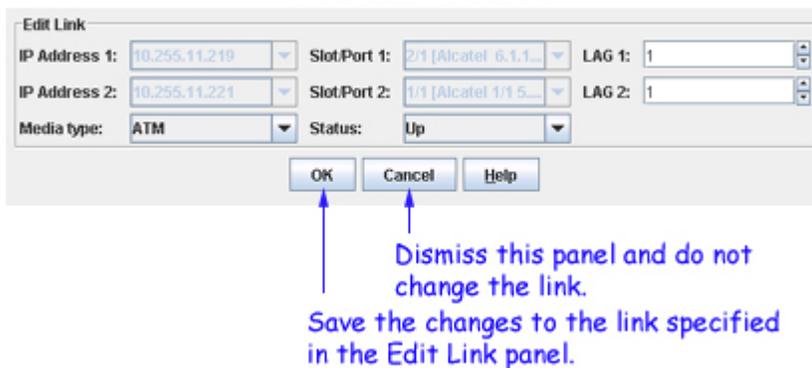
You can export the list of links from this application to a comma-separated value file (.csv file). This file can be displayed and edited in Microsoft Excel or any other application that uses comma-separated value files. To export a list of links, select **Export Manual Links** on the File menu. The Export Links window displays, as shown below. Locate the area where you want to save the link list file, enter a name for the file, and then click the **Save** button.



Editing Links

To edit an existing manual link, select the link in the link list and click the **Edit** button. The Edit Link panel activates, as shown below. When you have made the changes desired, click the **OK** button and then click the **Apply** button to save the change to the server.

The Edit Link Panel



Deleting Links

To delete a link, select the link and click the **Delete** button.

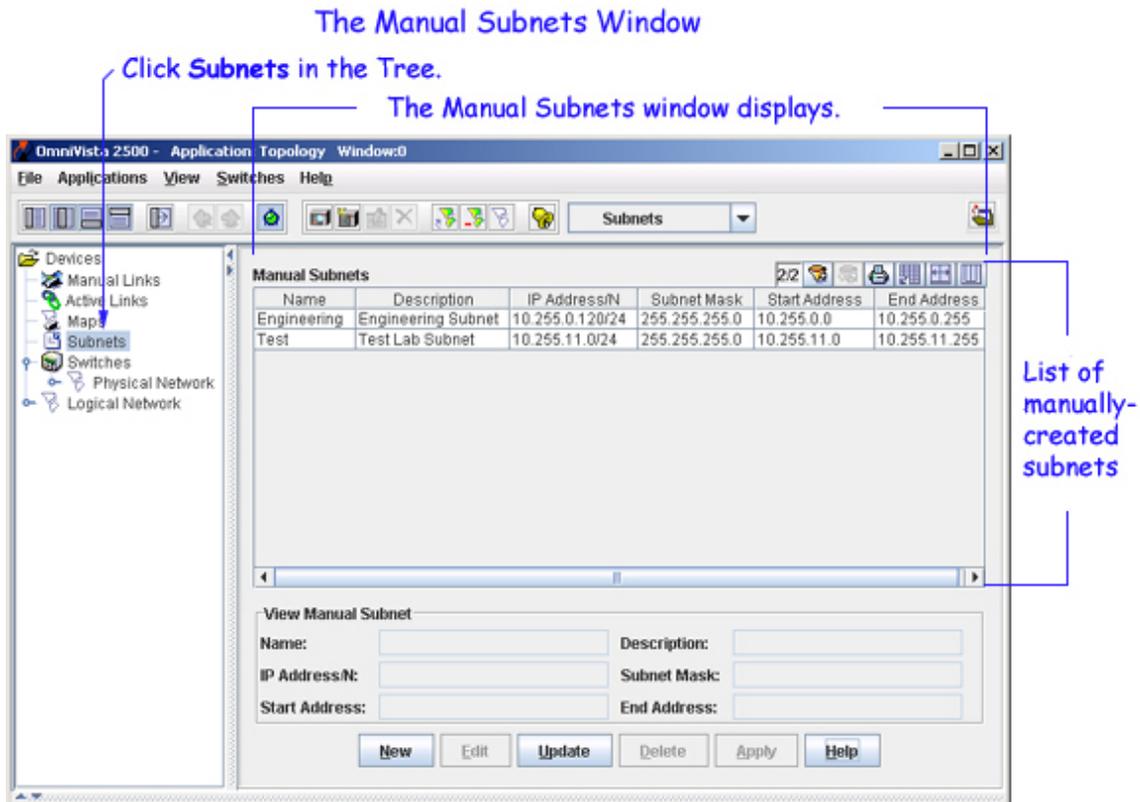


Managing Manual Subnets

The Manual Subnets window enables you to override OmniVista's default subnet creation and manually define the subnets that OmniVista displays in the tree. (OmniVista's default subnet creation is described below.) If manual subnets exist when a discovery is performed, OmniVista will place the discovered switches into the manual subnets upon discovery. If manual subnets are created after discovery, OmniVista will place known switches into the manual subnets upon their creation.

The subnets displayed in the tree will be updated and their contents adjusted automatically whenever the list of manual subnets is modified. OmniVista will automatically place each switch in the subnet that most specifically defines it. If a subnet is created but no known switch falls into the range of the subnet, that subnet will not be displayed in the tree. Manual subnets can be subsets or supersets of existing subnets. A manual subnet cannot duplicate any existing manual subnet.

Note: Any additions or modifications made to the list of manual subnets will apply to all users logged on to the current OmniVista server. For this reason only users with Admin or Net Admin security privileges are allowed to add, modify, or delete manual subnets.



Subnet Names

In the tree, subnets are labeled in the form *ipaddress/n*. The */n* indicates the number of bits in *ipaddress*, starting from the left, that identify the network (i.e., the subnet). These bits will have the same value in all the addresses that belong to the subnet. The literal value of these bits displays in *ipaddress*. Any bits in *ipaddress* that do not identify the subnet are represented by zeros.

For example, the screen above shows a subnet named **10.255.11.0/24**. The **/24** means that the first 24 bits of the address, starting from the left, identify the subnet and will be common to all address in the subnet. The literal value of these 24 bits, 10.255.11, displays in the subnet name. The last bits are represented by a 0, as these bits do not identify the subnet. (They identify devices.) This subnet could also be represented as 10.255.11.*, where the * character represents any value. This subnet will include all devices with an IP address in the range 10.255.11.0 - 10.255.11.255.

As a second example, consider a subnet named **10.0.0.0/8**. The **/8** means that the first eight bits of the address identify the subnet and will be common to all address in the subnet. The literal value of these eight bits, 10, displays in the subnet name. All other bits are represented by zeros. This subnet could also be represented as 10.*.*.*, where the * character represents any value. This subnet will include all the devices with an IP address in the range 10.0.0.0 - 10.255.255.255.

Default Subnet Creation

By default, OmniVista places the switches it discovers into subnets according to the Class of the switch IP addresses, as follows:

Class C addresses. IP addresses that start with a decimal value of 192 or higher -- such as 192.10.20.30 or 200.15.53.33 -- are assumed to belong to a Class C subnet, wherein the first three decimal values of the IP address, starting from the left, identify the subnet. For example, OmniVista would place IP address 200.15.53.33 in subnet 200.15.53.*. This subnet could also be represented as 200.15.53.0/24, where the "/24" means that the first 24 bits of the address, starting from the left, identify the network (i.e., the subnet) in which the address belongs.

Class B addresses. IP addresses that start with a decimal value between 127 and 191, inclusive -- such as 127.10.20.30 or 150.15.53.33 -- are assumed to belong to a Class B subnet, wherein the first two decimal values of the IP address, starting from the left, identify the subnet. For example, OmniVista would place IP address 150.15.53.33 into subnet 150.15.*.*. This subnet could also be represented as 150.15.0.0/16, where the "/16" means that the first 16 bits of the address, starting from the left, identify the network (i.e., the subnet) in which the address belongs.

Class A addresses. IP addresses that start with a decimal value of 126 or lower -- such as 10.10.20.30 or 120.15.53.33 -- are assumed to belong to a Class A subnet, wherein the first decimal value of the IP address, starting from the left, identifies the subnet. For example, OmniVista would place IP address 120.15.53.33 into subnet 120.*.*.*. This subnet could also be represented as 120.0.0.0/8, where the "/8" means that the first eight bits of the address, starting from the left, identify the network (i.e., the subnet) in which the address belongs.

Example: Creation of a Manual Subnet

As an example, let's say that after a discovery is performed all discovered switches are displayed in one default subnet, which is labeled 10.0.0.0/8 in the tree. The network administrator then creates manual subnet 10.255.11.0/24. When the administrator clicks **Apply** in the Manual Subnets window to create this new subnet, all devices that have an IP address within the range 10.255.11.0 - 10.255.11.255 will move into the new manual subnet. If no switches then remain in default subnet 10.0.0.0/8, that subnet will be removed from the Tree and will no longer display. If the administrator later deletes manual subnet 10.255.11.0/24 from the list of manual subnets, default subnet 10.0.0.0/8 will redisplay in the tree with all original member switches.

How to Create a Manual Subnet

Follow the steps below to create a new manual subnet.

1. Click the **New** button. The Create Manual Subnet fields are activated, as shown below. In the **Name** field, enter a short name that describes the subnet.

Note: This name will be used instead of the default "Subnet (ipaddress/n)" name in the tree, and elsewhere, when the switch name preferences is set to either "System Name" or "DNS Name" in the Preferences application. If the switch name preference is set to "IP Only", it will have no effect.

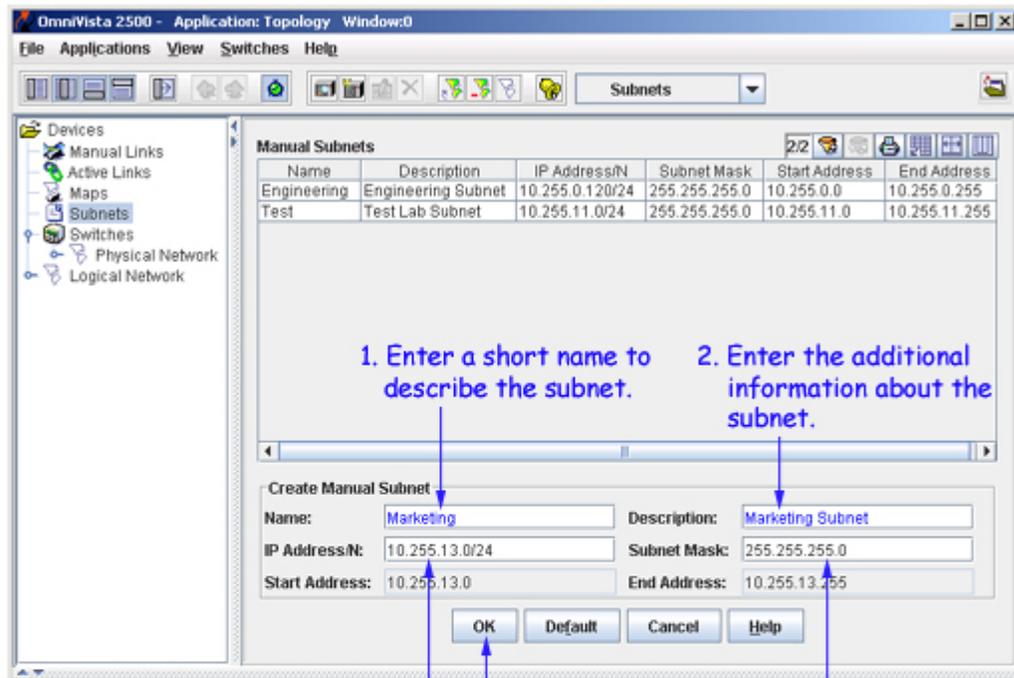
2. In the **Description** field, enter any additional information about the subnet.

3. In the **IP Address/N** field, enter any address that belongs in the new subnet, followed by /*n*, where *n* is the number of bits in the IP address that identify the subnet. These bits will be common to all addresses in the subnet.

4. In the **Subnet Mask** field, enter the mask for the new subnet. As soon as the subnet mask is entered, the start address and end address thus defined for the new subnet automatically display in the **Start Address** and **End Address** fields.

5. If the subnet is correctly defined, click the **OK** button to create the new manual subnet. Alternatively, you can enter new values in the **IP Address/N** and/or **Subnet Mask** fields to redefine the subnet until the desired start and end addresses display. The new subnet displays in the list of Manual Subnets after you click the **OK** button.

Creating a New Manual Subnet



1. Enter a short name to describe the subnet.
2. Enter the additional information about the subnet.
3. Enter the IP address of the new subnet, followed by /n, where n is the number of bits in the IP address that identify the subnet. These bits will be common to all addresses in the subnet.
4. Enter the subnet mask. The start and end addresses thus defined for the new subnet display in the **Start Address** and **End address** fields.
5. If the subnet is correctly defined, click the **OK** button to create the subnet.

6. Click the **Apply** button to write the new manual subnet to the server. When you click **Apply**, OmniVista populates the new subnet with all switches that fall within its range and reorders the Tree display accordingly.

Note: When you click the **Update** button, the list of all the manually created subnets in the OmniVista will be refreshed and displayed.

Editing a Manual Subnet

1. Select the subnet in the list of manual subnets and click the **Edit** button. The Edit Manual Subnet fields activate.
2. Edit the **Name**, **Description**, **IP Address/N**, or the **Subnet Mask** field as desired.
3. Click the **OK** button when your changes are complete. In the list of Manual Subnets, the subnet that you edited is marked as a pending deletion and a new subnet that reflects your changes is marked as a pending addition.

4. Click the **Apply** button to write the changes to the server. OmniVista deletes the previous subnet and adds the new subnet that reflects your changes. The tree display is updated and switches are reassigned to the subnets accordingly. If no appropriate manual subnet exists for a switch, OmniVista will create an appropriate default subnet as described above.

Deleting a Manual Subnet

1. Select the subnet in the list of Manual Subnets and click the **Delete** button.

2. Click the **Apply** button to write the change to the server. When you click **Apply**, OmniVista deletes the subnet and the tree display is updated accordingly. If no appropriate manual subnet exists for a switch, OmniVista will create an appropriate default subnet as described above.

Connecting to a Switch

You can connect to a switch merely by selecting it in the Tree. When the connection is established, information about the switch displays, as shown below. Note that the information displayed is somewhat different for AOS devices and XOS devices.

To connect to a switch, select it in the Tree

When the connection is established, switch information displays

The screenshot displays the OmniVista 2500 application interface. On the left, a tree view shows a 'Physical Network' with several IP addresses. The IP address '10.255.11.97' is highlighted, indicating it is selected. On the right, a configuration window for the device 'kite2_97 (10.255.11.97)' is open. The window has tabs for 'General', 'Modules', 'Physical Port', 'Interfaces', 'Statistics', 'Link Agg', 'Link Agg Ports', and 'System State'. The 'General' tab is active, showing system information such as Name, Description, Location, Contact, Up Time, and AMAP. Below this, there are configuration options for 'Running From', 'Current State', 'Certify State', 'Synchronization State', 'Save To Working', 'Load From', and 'Synchronize CMMs'. At the bottom, there is a 'Chassis Information Table' with columns for Primary CMM, Free Slots, Number of Resets, Board Temp (°C), Temp Range, Temp Threshold (°C), and Danger. The table contains one row of data for 'CMM in Slot 1'.

| Primary CMM | Free Slots | Number of Resets | Board Temp (°C) | Temp Range | Temp Threshold (°C) | Danger |
|---------------|------------|------------------|-----------------|----------------|---------------------|--------|
| CMM in Slot 1 | 7 | 20 | 37 | underThreshold | 57 | |

General Tab (AOS Devices)

The General tab for AOS devices provides general system information and specific chassis information. It also enables you to start and stop the AMAP protocol and to save, load, copy, and synchronize switch configuration files, as explained in detail below. Note that you can change any configurable parameter by editing the field desired and then clicking **Apply** to write the change to the switch. All changes take effect immediately.

The General Tab

Device: FALCON-137 (10.255.11.221)

General Modules Physical Port Interfaces Statistics Link Agg Link Agg Ports System State

System

Name: FALCON-137 Description: 5.4.1.167.R01 Development, January 11, 2006.
 Location: Unknown Contact: Alcatel Internetworking, www.alcatel.com/enterpr
 Up Time: 2 days 9 hours 7 minutes 30 seconds AMAP: Active

Configuration

Running From: Working - CMM-A Current State: Saved
 Certify State: Need Certify Synchronization State: Not Applicable
 Save To Working: Load From: Working
 Copy: Working To Certified Synchronize CMMs:

Chassis Information Table

| Primary CMM | Free Slots | Power Left(Watts) | Number of Resets | Board Temp.(°C) | CPU Temp.(°C) | Temp Range |
|-------------|------------|-------------------|------------------|-----------------|---------------|----------------|
| CMM-A | 15 | 279 | 3078 | 48 | 35 | underThreshold |

Apply Update Help

Click Update to poll the switch and refresh the screen with current information.

Click Apply to write changes to the switch. All changes take effect immediately.

System Parameters

System

Name: kite2_97 Description: 6.1.2.144.R01 Development, January 06, 2006
 Location: Unknown Contact: Alcatel Internetworking, www.alcatel.com/enter
 Up Time: 1 hours 20 minutes 30 seconds AMAP: Active

Name

A user-defined name for this switch.

Description

A factory-defined description of the switch's software.

Location

A user-defined description of the switch's physical location.

Contact

A user-defined statement identifying the person or organization responsible for the switch.

Up Time

The period of time that has elapsed since the switch was last rebooted.

AMAP

Set this field to **Active** or **Inactive** to enable or disable the AMAP protocol on this switch. By default, AMAP is enabled. AMAP is a proprietary protocol that learns the connections and links between switches in the list of All Discovered Devices. This information is used to create a graphical display of network links when a network region or subnet is viewed.

Saving and Loading Configuration Files

Overview

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 -- refer to the Load From Working and Load From Certified commands described below.)

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 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.

Configuration Parameters

Running From

This read-only field displays the directory and CMM module from which configuration files were originally loaded: either the **working** directory, the **certified** directory, or **unknown**. When the configuration files were loaded from the working directory, you are allowed to save configuration changes and the **Save To Working** checkbox is enabled. When the configuration files were loaded from the certified directory, you are not allowed to save configuration changes and the **Save To Working** checkbox is disabled.

Current State

This read-only field displays the current state of the CMM's running configuration: either **saved**, **unsaved**, or **uncertified**.

saved. The running configuration is identical to the contents of the directory from which the configuration files were originally loaded -- either the working directory or the certified directory.

unsaved. The running configuration has been changed and is not identical to the contents of the directory from which the configuration files were originally loaded.

uncertified. The working directory contains saved configuration changes that are not in the certified directory. The working directory and the certified directory are different.

Certify State

This read-only field reports the certification state of the CMM's working directory; that is, whether the working directory matches the certified directory.

Certified. The CMM's working directory is identical to its certified directory.

Need Certify. The CMM's working directory is not identical to its certified directory.

Unknown. The CMM's certification state is unknown.

Synchronization State

This read-only field reports 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.

Save To Working

This checkbox can be enabled only when the CMM is running from (i.e., originally loaded from) the working directory. It enables you to save the running configuration of the CMM to the working directory. If you save the configuration to the working directory, the **Current State** field, described above, will change to **uncertified**. Note that it may take up to 1 1/2 minutes for the **Current State** field to update.

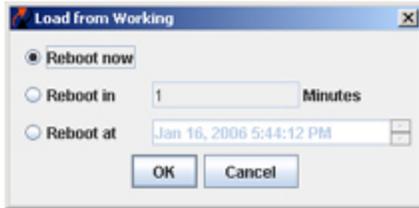
Note: When you apply **Save to Working** on a device, you must allow 120 seconds of time to elapse, before you apply the same again.

Load From Working

This checkbox enables you to reboot the primary CMM from the **Working** directory. Note that any unsaved configuration changes will be lost: you can save configuration changes with the **Save To Working** command before executing **Load From Working**.

When you select **Load From Working** and click **Apply**, the Load from Working window displays. The Load from Working window is shown below. This window enables you to specify whether you wish to reboot immediately (**Reboot now**), or reboot within 1 - 1000 minutes (**Reboot in x Minutes**), or reboot at a specified date and time (**Reboot at date time**). Specify the desired reboot time and then click the **OK** button.

The Load from Working window enables you to schedule the reboot.



Load From Certified

This checkbox enables you to reboot the primary CMM from the **Certified** directory. Note that any unsaved configuration changes will be lost: you can save configuration changes with the **Save To Working** command before executing **Load From Certified**.

When you select **Load From Certified** and click **Apply**, the Load from Certified window displays. The Load from Certified window is shown below. This window enables you to specify whether you wish to reload an entire switch (**Reload Entire Switch**), reboot immediately (**Reboot now**), or reboot within 1 - 1000 minutes (**Reboot in x Minutes**), or reboot at a specified date and time (**Reboot at date time**). Specify the desired reboot time and then click the **OK** button.

The Load from Certified window enables you to schedule the reboot.



Note: When you reboot the primary CMM from the certified directory, the switch will automatically failover to the secondary CMM (in other words, the two CMMs will trade primary and secondary roles). When you reboot the primary CMM from the working directory, no failover occurs. When rebooting from the certified directory, you should first synchronize the primary and secondary CMMs in order to ensure effective redundancy prior to failover.

Copy Certified to Working or Working to Certified

Depending on your selection, enabling this checkbox and clicking **Apply** causes the contents of the certified directory in the primary CMM to be copied to the working directory in the primary CMM, or causes the contents of the working directory in the primary CMM to be copied to the certified directory in the primary CMM.

Note: To prevent conflict between two long-running operations (such as, save to working, copy to working, etc.) on the same switch, OmniVista locks the conflicting running operations for a small duration of time.

Synchronize CMMs

Enabling this checkbox and clicking **Apply** causes the contents of the certified and the working directories in the primary CMM to be copied to the secondary CMM. By synchronizing the two CMM modules, the

switch has effective redundancy any time a failover occurs. It is recommended that you apply this function before reloading your primary CMM.

Chassis Information Parameters

| Chassis Information Table | | | | | | |
|---------------------------|------------|------------------|-----------------|----------------|---------------------|----------------------------|
| Primary CMM | Free Slots | Number of Resets | Board Temp (°C) | Temp Range | Temp Threshold (°C) | Danger Temp Threshold (°C) |
| CMM in Slot 1 | 7 | 21 | 37 | underThreshold | 57 | 94 |

Note: Not all fields display for all devices. If a field is not applicable to a device it is not displayed.

Primary CMM

This field identifies the CMM that is currently functioning as the primary CMM.

Free Slots

The number of Network Interface front panel slots that are empty.

Power Left (Watts)

The amount of power still available on the chassis, in Watts.

Number of Resets

The number of times this switch has been reset since the last cold start.

Board Temp (Degrees Celsius)

The current reading of the board temperature sensor, in degrees Celsius, for this chassis. The value in this field is compared to the **Temp Threshold** value (described below) for purposes of determining if the **Board Temp** is over or under the threshold value. The result of this comparison is displayed by the **Temp Range** parameter (described below).

CPU Temp (Degrees Celsius)

The current reading of the SPARC temperature sensor, in degrees Celsius, for this chassis.

Temp Range

This field displays the results of the comparison of the **Board Temp** value (described above) and the **Temp Threshold** value (described below). This field also indicates if the **Board Temp** value is over the **Danger Temp Threshold** value (described below). The value in this field can display as:

unknown. The comparison value is unknown.

not Present. A value required for the comparison is not present.

underThreshold. The Board Temp value is lower than the Temp Threshold value.

overFirstThreshold. The **Board Temp** value is higher than the **Temp Threshold** value but is lower than the **Danger Temp Threshold** value.

overDangerThreshold. The Board Temp value is higher than the Danger Temp Threshold value.

Temp Threshold (Degrees Celsius)

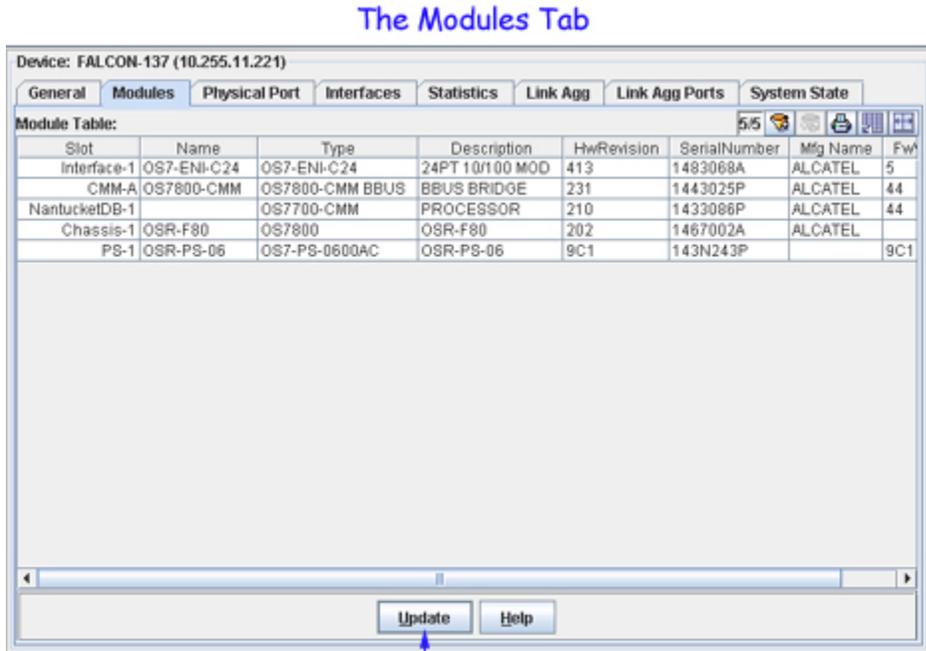
This threshold value, in degrees Celsius, is the temperature level at which -- when reached due to either an ascending or descending temperature transition -- temperature notification is provided to the user. When this threshold value is exceeded, traps and other operator notifications are transmitted.

Danger Temp Threshold (Degrees Celsius)

The Danger Temperature Threshold is factory-configured at 80 degrees Celsius and cannot be changed. If the chassis should exceed this temperature it will start shutting down Network Interface modules.

Modules Tab (AOS Devices)

The Modules tab lists the hardware modules installed in the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each column is described below.



Slot

The slot in which the module is installed.

Name

The name of the module

Type

The factory-defined physical type of the module.

Description

A description of the module.

HwRevision

The current revision level of the module hardware

SerialNumber

Serial number of the module.

Mfg Name

The name of the manufacturer.

FwVersion

The module's firmware version. All modules should use the same firmware version.

Physical Port Tab (AOS Devices)

The Physical Port tab provides information on all physical ports on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Physical Port Tab

Device: FALCON-137 (10.255.11.221)

General Modules **Physical Port** Interfaces Statistics Link Agg Link Agg Ports System State

BOP Physical Port Status Table 24/24

| Slot | Port | Media Type | Alias | Description | Admin. Status | Oper. Status |
|------|------|------------------|-------|----------------------------|---------------|--------------|
| 1 | 1 | ETHERNET-CSMA/CD | | Alcatel 1/1 5.4.1.167.R01 | portUp | down |
| 1 | 2 | ETHERNET-CSMA/CD | | Alcatel 1/2 5.4.1.167.R01 | portUp | down |
| 1 | 3 | ETHERNET-CSMA/CD | | Alcatel 1/3 5.4.1.167.R01 | portUp | down |
| 1 | 4 | ETHERNET-CSMA/CD | | Alcatel 1/4 5.4.1.167.R01 | portUp | up |
| 1 | 5 | ETHERNET-CSMA/CD | | Alcatel 1/5 5.4.1.167.R01 | portUp | down |
| 1 | 6 | ETHERNET-CSMA/CD | | Alcatel 1/6 5.4.1.167.R01 | portUp | down |
| 1 | 7 | ETHERNET-CSMA/CD | | Alcatel 1/7 5.4.1.167.R01 | portUp | down |
| 1 | 8 | ETHERNET-CSMA/CD | | Alcatel 1/8 5.4.1.167.R01 | portUp | down |
| 1 | 9 | ETHERNET-CSMA/CD | | Alcatel 1/9 5.4.1.167.R01 | portUp | down |
| 1 | 10 | ETHERNET-CSMA/CD | | Alcatel 1/10 5.4.1.167.R01 | portUp | down |
| 1 | 11 | ETHERNET-CSMA/CD | | Alcatel 1/11 5.4.1.167.R01 | portUp | down |
| 1 | 12 | ETHERNET-CSMA/CD | | Alcatel 1/12 5.4.1.167.R01 | portUp | down |
| 1 | 13 | ETHERNET-CSMA/CD | | Alcatel 1/13 5.4.1.167.R01 | portUp | down |
| 1 | 14 | ETHERNET-CSMA/CD | | Alcatel 1/14 5.4.1.167.R01 | portUp | down |
| 1 | 15 | ETHERNET-CSMA/CD | | Alcatel 1/15 5.4.1.167.R01 | portUp | down |
| 1 | 16 | ETHERNET-CSMA/CD | | Alcatel 1/16 5.4.1.167.R01 | portUp | down |
| 1 | 17 | ETHERNET-CSMA/CD | | Alcatel 1/17 5.4.1.167.R01 | portUp | down |
| 1 | 18 | ETHERNET-CSMA/CD | | Alcatel 1/18 5.4.1.167.R01 | portUp | down |
| 1 | 19 | ETHERNET-CSMA/CD | | Alcatel 1/19 5.4.1.167.R01 | portUp | down |
| 1 | 20 | ETHERNET-CSMA/CD | | Alcatel 1/20 5.4.1.167.R01 | portUp | down |
| 1 | 21 | ETHERNET-CSMA/CD | | Alcatel 1/21 5.4.1.167.R01 | portUp | down |

Update Help

Click Update to poll the switch and refresh the screen with current information.

Slot/Port

The slot and port for which status is displayed.

Media Type

The physical type of the port.

Alias

The user-defined alias for the port.

Description

A description of the port.

Admin Status

The Administrative (Admin) status of the port: **up** or **down**. When the Admin status of a port is enabled, the port can receive and transmit data as long as a cable is connected and no physical or operational problems exist. When the Administrative Status of a port is disabled, 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 Administrative Status is enabled.

OperStatus

The operational status of the port: **portUp**, **portDown**, or **unknown**.

The Interfaces Tab (AOS Devices)

The Interfaces tab provides information about each physical interface in the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Interfaces Tab

| Index | Alias | Description | Type | MTU | Speed | Physical Address | Admin. State |
|-------|--------------|---------------|------------------|------|---------|-------------------|--------------|
| 1001 | Alcatel 1/1 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 c4 | up |
| 1002 | Alcatel 1/2 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 c5 | up |
| 1003 | Alcatel 1/3 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 c6 | up |
| 1004 | Alcatel 1/4 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 100 Mbs | 00 d0 95 6b 26 c7 | up |
| 1005 | Alcatel 1/5 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 c8 | up |
| 1006 | Alcatel 1/6 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 c9 | up |
| 1007 | Alcatel 1/7 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 ca | up |
| 1008 | Alcatel 1/8 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 cb | up |
| 1009 | Alcatel 1/9 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 cc | up |
| 1010 | Alcatel 1/10 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 cd | up |
| 1011 | Alcatel 1/11 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 ce | up |
| 1012 | Alcatel 1/12 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 cf | up |
| 1013 | Alcatel 1/13 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 d0 | up |
| 1014 | Alcatel 1/14 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 d1 | up |
| 1015 | Alcatel 1/15 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 d2 | up |
| 1016 | Alcatel 1/16 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 d3 | up |
| 1017 | Alcatel 1/17 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 d4 | up |
| 1018 | Alcatel 1/18 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 d5 | up |
| 1019 | Alcatel 1/19 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 d6 | up |
| 1020 | Alcatel 1/20 | 5.4.1.167.R01 | ETHERNET-CSPA/CD | 1553 | 0 | 00 d0 95 6b 26 d7 | up |

Click Update to poll the switch and refresh the screen with current information.

Index

A unique value that identifies this interface internally.

Description

A description of the interface that usually includes the name of the manufacturer, the name of the product, and the version of the interface's hardware/software.

Type

A description of the type of the interface.

MTU

The size, in octets, of the largest packet that can be sent or received on the interface.

Speed

An estimate of the interface's current bandwidth. Speed is displayed in bits-per-second if less than 1,000,000 bits-per-second. Speeds of 1,000,000 bits-per-second or greater are displayed in terms of Mbs (megabits-per-second). If an interface does not vary in bandwidth, or if no accurate estimation can be made, the nominal bandwidth is displayed in this field.

Physical Address

The physical address of the interface at its protocol sublayer. For 802.x interfaces, the physical address is a MAC address. No physical address displays for interfaces in loopback mode nor for serial interfaces.

Admin. State

The administrative state of the interface: **up**, **down**, or **testing**. Admin state **up** indicates the interface is administratively enabled to pass packets; **down** indicates the interface is administratively disabled from passing packets; **testing** indicates the interface is in a test mode and cannot pass operational packets. All interfaces are initialized with the admin state **down**. After initialization, either in response to explicit management action or stored configuration data, the admin state of an interface to changed to **up** or **testing** (or may remain **down**).

Oper. Status

The current operational status of the interface: **up**, **down**, **testing**, **unknown**, **dormant**, **notPresent**, or **lowerLayerDown**.

up. The interface is ready to transmit and receive packets.

down. The interface is either administratively disabled or there is a fault that prevents it from going to the **up** state.

testing. The interface is in a test mode and cannot pass operational packets.

dormant. The interface is waiting for external actions (such as a serial line waiting for an incoming connection).

notPresent. The interface has missing components (typically hardware components).

lowerLayerDown. The interface is down due to the state of lower-layer interfaces.

If an interface's administrative state is **down** its operational status will also be **down**. When the administrative state 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 going **up**; or, the operational status will remain

Last Change

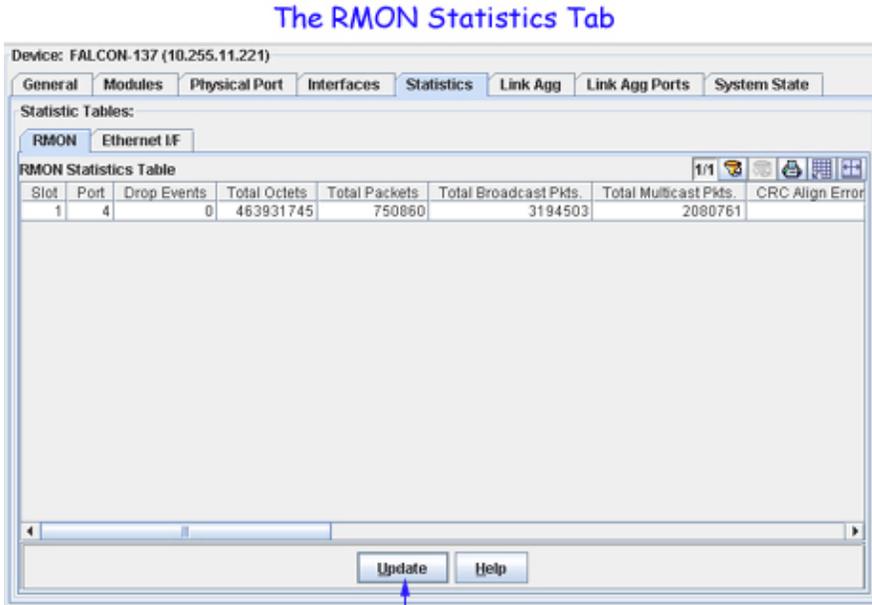
The value of sysUpTime when the interfaces table (ifTable) was last changed because a new entry was created or an existing entry was deleted. (The sysUpTime MIB variable reports the time period that has elapsed since the switch was last initialized.) If the interfaces table was not changed since the last reinitialization of OmniVista, no value will display in this field.

Out Queue

The length of the packet output queue, in packets.

RMON Statistics (AOS Devices)

The RMON Statistics tab, shown below, displays statistics for RMON (Remote Monitoring). Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.



Click **Update** to poll the switch and refresh the screen with current information.

Slot and Port

The slot and port for which RMON statistics are displayed.

Drop Events

The total number of occasions that packets were dropped by the probe due to lack of resources. Note that the value in this field is not necessarily the number of packets dropped; it is the number of times this condition was detected.

Total Octets

The total number of octets received, including those in bad packets. The count includes FCS (frame check sequence) octets but excludes framing bits. The value in this field can be used as a reasonable estimate of 10 megabit Ethernet utilization. If greater precision is desired, the **Total Octets** and **Total Packets** values should be sampled before and after a common interval. In the following equation, the differences in the sampled values are *Octets* and *Pkts*, respectively, and the number of seconds in the common interval is *Interval*. The result of this equation is the value *Utilization* which is the percent utilization of the Ethernet segment on a scale of 0 to 100 percent.

$$Utilization = \frac{Pkts * (9.6 + 6.4) + (Octets * .8)}{Interval * 10,000}$$

Total Packets

The total number of packets received, including bad packets, broadcast packets, and multicast packets.

Total Broadcast Pkts

The total number of good packets received that were directed to the broadcast address. Note that this value does not include multicast packets.

Total Multicast Pkts

The total number of good packets received that were directed to a multicast address. Note that this value does not include packets directed to the broadcast address.

CRC Align Errors

The total number of packets received with a length between 64 and 1518 octets, inclusive (excluding framing bits but including FCS [frame check sequence] octets), which had either of the following errors:

a bad frame check sequence with an integral number of octets, which is an FCS error, or

a bad frame check sequence with a non-integral number of octets, which is an alignment error.

Undersized Pkts

The total number of packets received that were less than 64 octets in length, excluding framing bits but including FCS (frame check sequence) octets, and were otherwise well formed.

Oversized Pkts

The total number of packets received that were longer than 1518 octets, excluding framing bits but including FCS (frame check sequence) octets, and were otherwise well formed.

Fragments

The total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS [frame check sequence] octets), which had either of the following errors:

a bad frame check sequence with an integral number of octets, which is an FCS error, or

a bad frame check sequence with a non-integral number of octets, which is an alignment error.

Note that it is entirely normal for the count in this field to increment, because it includes both runt packets (which are a normal occurrence due to collisions) and noise hits.

Jabbers

The total number of packets received that were longer than 1518 octets (excluding framing bits but including FCS [frame check sequence] octets), which had either of the following errors:

a bad frame check sequence with an integral number of octets, which is an FCS error, or

a bad frame check sequence with a non-integral number of octets, which is an alignment error.

Note that this definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.

Collisions

The best estimate of the total number of collisions on this Ethernet segment. The value returned depends on the location of the RMON probe. Section 8.2.1.3 (10BASE-5) and section 10.3.1.3 (10BASE-2) of IEEE standard 802.3 states that a station, when in receive mode, must detect a collision if three or more stations are transmitting simultaneously. A repeater port must detect a collision when two or more stations are transmitting simultaneously. Thus, a probe placed on a repeater port could record more collisions than would a probe connected to a station on the same segment.

Probe location plays a much smaller role when considering 10BASE-T. Section 14.2.1.4 (10BASE-T) of IEEE standard 802.3 defines a collision as the simultaneous presence of signals on the DO and RD circuits (transmitting and receiving at the same time). A 10BASE-T station can only detect collisions when it is transmitting. Thus, a probe placed on a station and a probe placed on a repeater should report the same number of collisions.

Note that an RMON probe inside a repeater should ideally report collisions between the repeater and one or more other hosts (per the IEEE 802.3k definition of transmit collisions) plus receiver collisions observed on any coax segments to which the repeater is connected.

Pkts 64 Octets

The total number of packets received, including bad packets, that were 64 octets in length. The count includes FCS (frame check sequence) octets but excludes framing bits.

Pkts 65-127 Octets

The total number of packets received, including bad packets, that were between 65 and 127 octets in length, inclusive. The count includes FCS (frame check sequence) octets but excludes framing bits.

Pkts 128-255 Octets

The total number of packets received, including bad packets, that were between 128 and 255 octets in length, inclusive. The count includes FCS (frame check sequence) octets but excludes framing bits.

Pkts 256-511 Octets

The total number of packets received, including bad packets, that were between 256 and 511 octets in length, inclusive. The count includes FCS (frame check sequence) octets but excludes framing bits.

Pkts 512-1023 Octets

The total number of packets received, including bad packets, that were between 512 and 1023 octets in length, inclusive. The count includes FCS (frame check sequence) octets but excludes framing bits.

Pkts 1024-1518 Octets

The total number of packets received, including bad packets, that were between 1024 and 1518 octets in length, inclusive. The count includes FCS (frame check sequence) octets but excludes framing bits.

Ethernet Statistics (AOS Devices)

The Ethernet I/F tab lists statistics for each Ethernet interface in the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below. Note that discontinuities can occur in statistics values upon reinitialization of the system.

The Ethernet I/F Statistics Tab

Device: FALCON-137 (10.255.11.221)

General Modules Physical Port Interfaces **Statistics** Link Agg Link Agg Ports System State

Statistic Tables:

RMON Ethernet I/F

Ethernet Interface Statistics Table 24/24

| Slot | Port | Index | Type | Rx Octets | Tx Octets | Rx Unicast Pkts. | Tx Unicast Pkts. | Rx I/F Di |
|------|------|-------|------------------|-----------|-----------|------------------|------------------|-----------|
| 1 | 1 | 1001 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 2 | 1002 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 3 | 1003 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 4 | 1004 | ETHERNET-CRMA/CD | 464321417 | 195262991 | 752659 | 571777 | |
| 1 | 5 | 1005 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 6 | 1006 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 7 | 1007 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 8 | 1008 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 9 | 1009 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 10 | 1010 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 11 | 1011 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 12 | 1012 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 13 | 1013 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 14 | 1014 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 15 | 1015 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 16 | 1016 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |
| 1 | 17 | 1017 | ETHERNET-CRMA/CD | 0 | 0 | 0 | 0 | |

Update Help

Click Update to poll the switch and refresh the screen with current information.

Slot and Port

The slot and port of the interface.

Index

A unique value that identifies the interface internally.

Type

The type of the interface.

Rx Octets

The total number of octets received on the interface, including framing characters.

Tx Octets

The total number of octets transmitted out of the interface, including framing characters.

Rx Unicast Pkts

The total number of unicast packets received on this interface and delivered to a higher layer. This value does not include packets addressed to a multicast or broadcast address.

Tx Unicast Pkts

The total number of unicast packets that higher-level protocols requested be transmitted from this interface, including packets that were discarded or not sent. This value does not include packets addressed to a multicast or broadcast address at this sublayer.

Rx I/F Discards

The number of received packets that were discarded even though no errors were detected in the packets that would have prevented them from being delivered to a higher-layer protocol. One possible reason for discarding such packets would be the need to free buffer space.

Tx I/F Discards

The number of outbound packets that were discarded even though no errors were detected in the packets that would have prevented them from being transmitted. One possible reason for discarding such packets would be the need to free buffer space.

Rx I/F Errors

The number of received packets that contained errors preventing them from being delivered to a higher-layer protocol.

Tx I/F Errors

The number of outbound packets that could not be transmitted because of errors.

Unknowns

The number of received packets that were discarded because of an unknown or unsupported protocol.

Link Agg Tab (AOS Devices)

Link aggregation is a way of combining multiple physical links between two switches into one logical link. The aggregate group operates within Spanning Tree as one virtual port and can provide more bandwidth than a single link. It also provides redundancy. If one physical link in the aggregate group goes down, link integrity is maintained.

There are two types of aggregate groups: static and dynamic. Static aggregate groups are manually configured on the switch with static links. Dynamic groups are set up on the switch but they aggregate links as necessary according to the Link Aggregation Control Protocol (LACP). OmniVista's Link Agg tab provides information about each link aggregation group defined on the switch. Each field in the tab is described below.



Click Update to poll the switch and refresh the screen with current information.

Number

A reference number assigned when the link aggregation group was created. This is a unique integer in the range 0 - 31 on OmniSwitch 6800/6850/7700/7800/9000 switches, 0 - 29 on OmniSwitch 6600 Family switches, and 0 - 15 on OmniSwitch 8800 switches.

Size

The maximum number of links that may belong to this link aggregation group. Values may be 2, 4, or 8, on an OmniSwitch 6800/6850/9000 switch. Values may be 2, 4, 8, or 16 on an OmniSwitch 7700, 7800, or 8800 switch. Values may 2, 4, or 8 on individual OmniSwitch 6600 Family switches and 2, 4, 8, or 16 on stacks consisting of two to eight OmniSwitch 6600 Family switches.

Name

The name of the link aggregation group. This is an alphanumeric string up to 255 characters long.

Description

The standard MIB name for this link aggregate group.

LACP Type

The type of this link aggregation group. **lacpOff** means the group is static. **lacpOn** means the group is dynamic and is using the LACP protocol. (LACP is the Link Aggregation Control Protocol.)

Admin State

The administrative state of this link aggregation group: either **enable** (the group is active and is able to aggregate links) or **disable** (the group is inactive). The group's administrative state is configured by the network administrator.

Oper State

The current operational state of this link aggregation group: either **up** (the group is operational) or **down** (the group is not operational). This field may also display **logicPortCreatFailed** or **qReservationFailed**.

Selected Ports

The number of ports that could possibly attach to this link aggregation group at the moment.

Attached Ports

The number of ports actually attached to this link aggregation group at the moment.

Primary Port

The slot/port number of the primary port in the link aggregation group used to send BPDUs and flooding frames. The switch uses the first port to join the group as the primary port. If the first port to join the group is no longer part of the group, the switch automatically assigns another port in the group to be the primary port.

MAC Address

The MAC address assigned to this link aggregation group.

Actor System ID

The MAC address for the local port associated with a dynamic link aggregation group, which is used as a unique identifier for the system that contains this link aggregation group.

Actor System Priority

A value from 0 - 65535 that indicates the priority value associated with the Actor System ID. This defines the priority of the switch's dynamic aggregate group in relation to other aggregate groups

Actor Admin Key

The administrative key value configured for the dynamic aggregate group. Possible values are 0 - 65535.

Actor Oper Key

The current operational value of the key for the dynamic link aggregation group.

Partner System ID

The MAC address of the remote aggregate group to which this aggregate group is attached. A value of zero indicates that there is no known partner. If the group is manually configured, the value in this field is assigned by the local system.

Partner System Priority

The priority of the remote system to which the aggregation group is attached. Possible values are 0 - 65535. If the group is manually configured, the value in this field is assigned by the local system.

Partner Admin Key

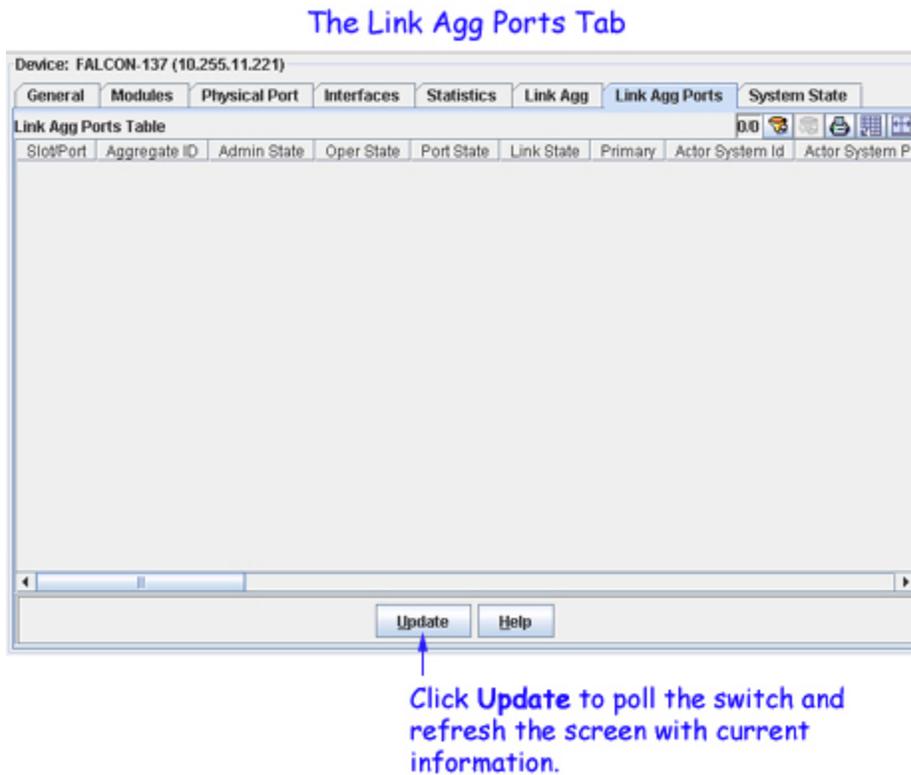
The administrative key for the aggregation group's remote partner. Possible values are 0 - 65535. If the group is manually configured, the value in this field is assigned by the local system. The administrative key may differ from the operational key.

Partner Oper key

The operational key of the remote system to which the aggregation group is attached. If the group is manually configured, the value in this field is assigned by the local system.

Link Agg Ports Tab (AOS Devices)

The Link Agg Ports tab provides information about the ports in link aggregation groups. Each field is described below.



Slot/Port

The slot and port number of a port in the link aggregation group.

Aggregate ID

The ID of the static aggregate group to which the port is attached. This field does not apply to dynamic aggregate groups. The **Aggregate ID** can be any value from **-1** to **31**. The **-1** value displays when this field is not significant.

Admin State

The administrative state of this port: either **enable** (the port is ready to pass packets) or **disable** (the port is administratively disabled). The port's administrative state is configured by the network administrator.

Oper State

The operational status of the port: either **up** (the port is passing traffic), **down** (the port is unable to pass traffic) **notAttached** (the port is not attached to the aggregate group), or **notAggregable** (the port cannot be aggregated, perhaps because the key is not set or is incorrect).

Port State

The current aggregation status of the port. When a port is attached to a group, **attached** will display in this field. Other possible port states are **created**, **configurable**, **configured**, **selected**, and **reserved**.

Link State

The operational status of the link: **up** or **down**.

Primary

This field displays **yes** if the port is the primary port in the aggregate group and displays **no** if it is not. This field may also display **notSignificant**.

Actor System ID

The System ID (i.e., the MAC address) of the system that contains this port.

Actor System Priority

A value from **0 - 255** that defines the priority value associated with the Actor's System ID.

Actor Admin Key

The actor administrative key value for this port.

Actor Oper Key

The current operational value of the actor key.

Partner Admin System ID

The administrative MAC address associated with the remote partner's system ID. This value is used along with Partner Admin System Priority, Partner Admin Key, Partner Admin Port, and Partner Admin Port Priority to manually configure aggregation.

Partner Oper System Priority

The operational priority of the remote system to which this port is attached.

Partner Admin Key

The administrative value of the key for the remote partner. This value is used along with Partner Admin System Priority, Partner Admin System ID, Partner Admin Port, and Partner Admin Port Priority to manually configure aggregation.

Partner Oper Key

The current operational value of the key for the protocol partner.

Selected Agg ID

The Aggregator ID associated with the dynamic aggregate group to which the port is attached. Zero indicates that this port has not selected an aggregate group, either because it is in the process of detaching from a group or because there is no suitable group available for it to select.

Attach Agg ID

The Aggregator ID associated with the dynamic aggregate group to which the port is attached. Zero indicates that this port is not currently attached to a group.

Actor Port

The port number locally assigned to this port. The port number is communicated in Link Aggregation Control Protocol Data Units (LACPDU) as the Actor_Port (a read-only value).

Actor Port Priority

The actor priority value assigned to the port. The actor priority value can range from 0 - 255.

Partner Admin Port

The administrative value of the port number for the protocol partner. This value is used along with Partner

Admin System Priority, Partner Admin System ID, Partner Admin Key, and Partner Admin Port Priority to manually configure aggregation.

Partner Oper Port

The operational port number assigned to the port by the port's protocol partner.

Partner Admin Port Priority

The administrative port priority of the protocol partner. This value is used along with Partner Admin System Priority, Partner Admin System ID, Partner Admin Key, and Partner Admin Port to manually configure aggregation.

Partner Oper Port Priority

The priority value assigned to this port by the partner.

Actor Admin State

The administrative state of the port. The Actor Admin State is a string of eight bits that correspond to the administrative values of Actor_State, as transmitted by the Actor in Link Aggregation Control Protocol Data Units (LACPDU)s. The bits of Actor Admin State are as follows:

The first bit corresponds to bit 0 of Actor_State, which is Activity. When this bit is set, the dynamic aggregate group is able to exchange LACPDU frames.

The second bit corresponds to bit 1 of Actor_State, which is Timeout. When this bit is set, a short timeout is used for LACPDU frames. When this bit is disabled, a long timeout is used for LACPDU frames.

The third bit corresponds to bit 2 of Actor_State, which is Aggregation. When this bit is set, the system considers this port to be a potential candidate for aggregation. If this bit is not enabled, the system considers the port to be individual (it can only operate as a single link).

The fourth bit corresponds to bit 3 of Actor_State, which is Synchronization. The system always determines the value of this bit. When bit 3 is set by the system, the port is allocated to the correct dynamic aggregation group. If this bit is not set by the system, the port is not allocated to the correct dynamic aggregation group.

The fifth bit corresponds to bit 4 of Actor_State, which is Collecting. The system always determines the value of this bit. When bit 4 is set by the system, incoming LACPDU frames are collected from the individual ports that make up the dynamic aggregate group.

The sixth bit corresponds to bit 5 of Actor_State, which is Distributing. The system always determines the value of this bit. When bit 5 is set by the system, distributing outgoing frames on the port is disabled.

The seventh bit corresponds to bit 6 of Actor_State, which is Defaulted. The system always determines the value of this bit. When bit 6 is set by the system, it indicates that the actor is using defaulted partner information administratively configured for the partner.

The eighth bit corresponds to bit 7 of Actor_State, which is Expired. The system always determines the value of this bit. When bit 7 is set by the system, the actor cannot receive LACPDU frames.

Actor Oper State

The operational state of the port. The Actor Oper State is a string of eight bits that correspond to the operational values of Actor_State, as transmitted by the Actor in Link Aggregation Control Protocol Data Units (LACPDU). The bits are allocated as described for **Actor Admin State** (see above).

Partner Admin State

The administrative state of the partner's port. The Partner Admin State is a string of eight bits that correspond to the administrative value of Actor_State for the protocol Partner.

The first bit corresponds to bit 0 of Actor_State for the Partner, which is Activity. When this bit is set, the dynamic aggregate group is able to exchange LACPDU frames.

The second bit corresponds to bit 1 of Actor_State for the Partner, which is Timeout. When this bit is set, a short timeout is used for LACPDU frames. When this bit is disabled, a long timeout is used for LACPDU frames.

The third bit corresponds to bit 2 of Actor_State for the Partner, which is Aggregation. When this bit is set, the system considers this port to be a potential candidate for aggregation. If this bit is not enabled, the system considers the port to be individual (it can only operate as a single link).

The fourth bit corresponds to bit 3 of Actor_State for the Partner, which is Synchronization. When this bit is set, the port is allocated to the correct dynamic aggregation group. If this bit is not enabled, the port is not allocated to the correct aggregation group.

The fifth bit corresponds to bit 4 of Actor_State for the Partner, which is Collecting. The system always determines the value of this bit. When bit 4 is set by the system, incoming LACPDU frames are collected from the individual ports that make up the dynamic aggregate group.

The sixth bit corresponds to bit 5 of Actor_State for the Partner, which is Distributing. The system always determines the value of this bit. When bit 5 is set by the system, distributing outgoing frames on the port is disabled.

The seventh bit corresponds to bit 6 of Actor_State for the Partner, which is Defaulted. The system always determines the value of this bit. When bit 6 is set by the system, it indicates that the partner is using defaulted actor information administratively configured for the actor.

The eighth bit corresponds to bit 7 of Actor_State for the Partner, which is Expired. The system always determines the value of this bit. When bit 7 is set by the system, the partner cannot receive LACPDU frames.

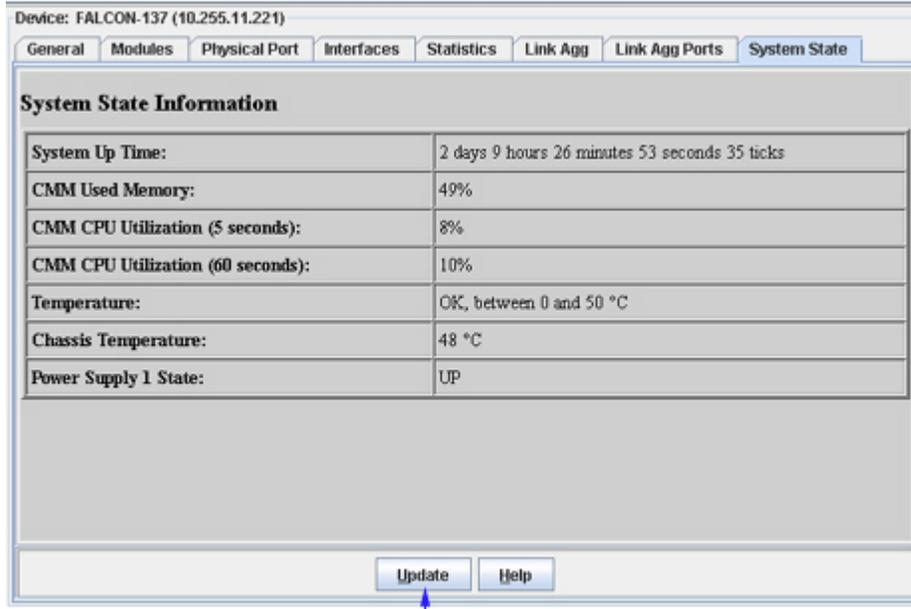
Partner Oper State

The current operational state of the partner's port. The Partner Oper State is a string of eight bits that correspond to the current values of Actor_State in the most recently received Link Aggregation Control Protocol Data Unit (LACPDU) transmitted by the protocol Partner. The bits are allocated as described for **Partner Admin State** (see above).

System State Tab (AOS Devices)

The System State tab, shown below, displays parameters that report the system state of the switch. Each field is described below.

The System State Tab



Click **Update** to poll the switch and refresh the screen with current information.

System Up Time

The time period that has elapsed since the switch was last initialized. (Each tick is .01 second.)

CMM Used Memory

The average device-level memory utilization, expressed as a percent, in the primary (active) CMM module over the latest sampling period.

CMM CPU Utilization (5 seconds)

The average device-level CPU utilization, expressed as a percent, in the primary (active) CMM module over the latest sampling period (every five seconds).

CMM CPU Utilization (60 seconds)

The average device-level CPU utilization, expressed as a percent, in the primary (active) CMM module over the last 60 seconds.

Temperature

This field indicates whether the chassis temperature is within the acceptable temperature range for the switch.

Chassis Temperature

The actual average temperature of the switch chassis, in degrees Celsius, over the latest sampling period.

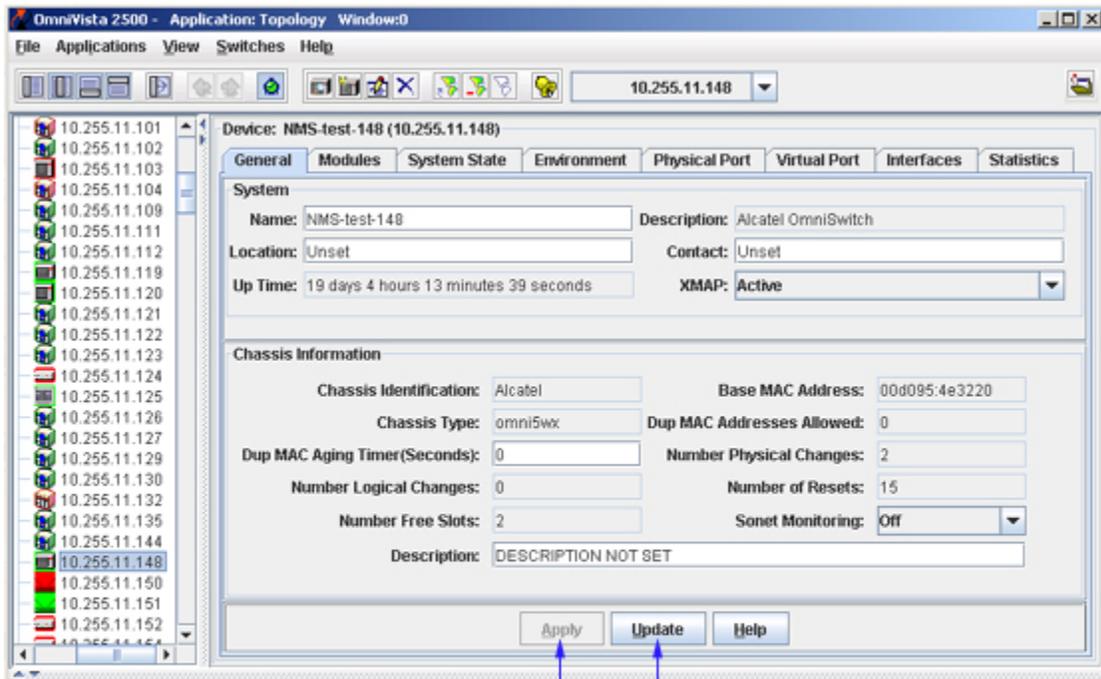
Power Supply *x* State

One instance of this parameter displays for each power supply that is present in the switch (for example, **Power Supply 1 State**, **Power Supply 2 State**, etc). If the maximum of four power supplies is present, four instances of this parameter will display. The power supply state is reported as **up** (the power supply is functional) or **down** (the power supply is not functional).

General Tab (XOS Devices)

The General tab for XOS devices provides general system information and chassis information, as explained in detail below. It also enables you to start and stop the XMAP protocol. To change any parameter, edit the field as desired and then click **Apply** to write the change to the switch. All changes take effect immediately.

The General Tab



Click **Update** to poll the switch and refresh the screen with current information.

Click **Apply** to write changes to the switch. All changes take effect immediately.

System Information Parameters

Name

A user-defined name for this switch.

Description

A description of the switch as defined by the manufacturer.

Location

A user-defined description of the switch's physical location.

Contact

A user-defined parameter stating who is responsible for this switch.

Up Time

The period of time that has elapsed since the switch was last rebooted.

XMAP

Set this field to **Active** or **Inactive** to enable or disable the XMAP protocol on this switch. By default, XMAP is enabled. XMAP is a proprietary protocol that learns the connections and links between switches in the list of Discovered Devices. This information is used to create a graphical display of network links when a network region or subnet is viewed. If you disable XMAP, this switch's connections and links will not be displayed.

Chassis Information Parameters

Note: Not all fields display for all devices. If a field is not applicable to a device it is not displayed.

Chassis Identification

This field identifies the manufacturer of the device being managed.

Base MAC Address

The base MAC address for the chassis is the first MAC address stored in the MPM. All MAC addresses associated with the MPM are derived from this base MAC address.

Chassis Type

The type of the chassis.

Dup MAC Addresses Allowed

Number of duplicate MAC addresses allowed on the switch.

Dup MAC Aging Timer (Seconds)

This field can be set to any value from 0 - 1000000 seconds. When set to a non-zero value, the Dup MAC Aging Timer specifies the aging time, in seconds, for duplicate MAC addresses learned from any Group in the switch. When set to zero, this timer is ignored and the Bridge Forwarding Table Aging Time value for the Group where the address was learned is used instead. Enabling the Dup MAC Aging Timer enables you to specify a chassis-wide aging time for duplicate MAC addresses.

Number Physical Changes

The number of physical changes that have been made to the switch since it was last reset or powered on. This includes the addition or removal of modules and controllers.

Number Logical Changes

The number of logical changes that have been made to the switch since it was last reset or powered on. This includes all sets to name strings.

Number of Resets

The number of times this switch has been reset since it was last cold-started.

Number Free Slots

The number of empty front-panel slots in the chassis.

SONET Monitoring

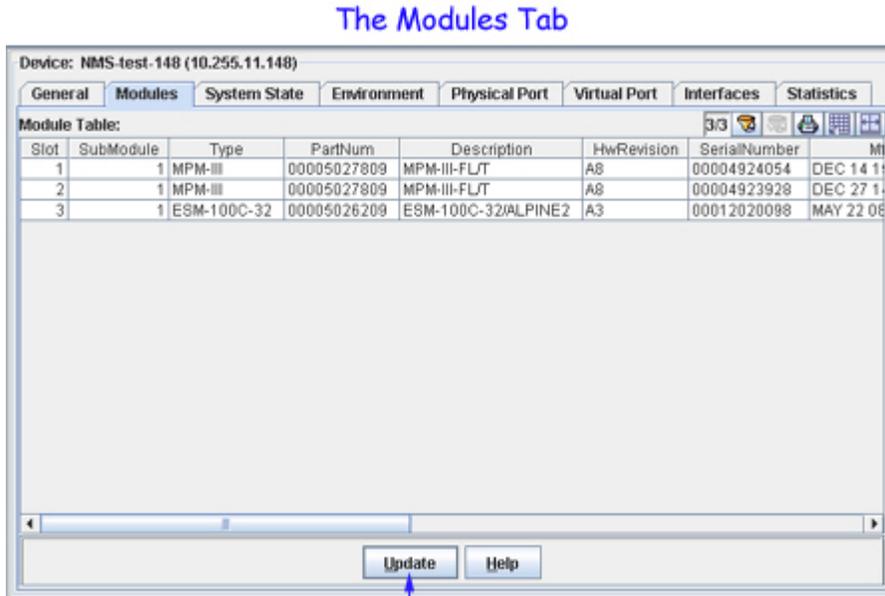
Set this field to **On** or **Off** to enable or disable SONET monitoring. The default value is Off. When this field is enabled, SONET error statistics are collected. Any change to this field takes effect as soon as the **Apply** button is clicked. (You can click **Update** to refresh the screen and see the new setting.) The SONET monitoring state applies to all CSM or ASM/ASX ports on the switch.

Description

A user-defined description of the switch chassis.

Modules Tab (XOS Devices)

The Modules tab lists the hardware modules installed in the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each column is described below.



Slot

The slot in which the module is installed. A switch chassis consists of numbered slots, which house various modules. Stackable switch models also have virtual slots and modules.

SubModule

Identifies the base module and any submodules present in a slot. The following identification scheme is used:

- 1 - base module
- 2 - submodule installed in the first position of the base module
- 3 - submodule installed in the second position of the base module
- 4 - submodule installed in the third position of the base module

Type

The physical type of the base module or submodule.

PartNum

The factory-assigned part number.

Description

A description of the module or submodule.

HwRevision

The current revision level of the module or submodule hardware.

SerialNumber

Serial number of the module or submodule.

MfgDate

The manufacturing date of the module or submodule.

FwVersion

The module or submodule's firmware version. All modules should use the same firmware version.

MAC Address

The base MAC address for this module or submodule. If the module or submodule does not support MAC addresses, the value in this field will be all zeros.

TimeStamp

The value of the sysUpTime MIB variable at the time this module was last reset.

AdminStatus

The administrative status of the module or submodule. Possible values are: Invalid, Enable, Disable, Reset, Load, Test, or Unknown (none of the previous).

OperStatus

The operational status of the module or submodule: Operational, Disabled, or Unknown ("Unknown" means uninitialized or that the module is in a transitional state).

VbusTxDiscards

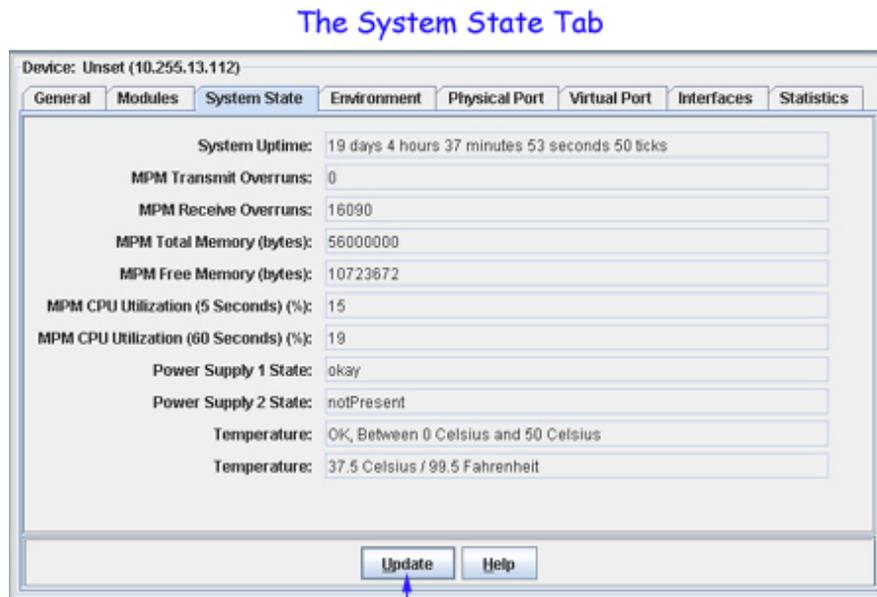
The current count of transmit VBUS buffer overruns.

VbusRxDiscards

The current count of receive VBUS buffer overruns.

System State Tab (XOS Devices)

The System State tab provides information related to the overall system state, system power, and chassis environment. All fields are read-only. Each field is described below.



System Uptime

The time since the last boot that the system has been running, displayed in days, hours, minutes, seconds, and ticks. (A tick is .01 second.)

MPM Transmit Overruns

The number of times a VSE transmit buffer could not be allocated by a task on the MPM.

MPM Receive Overruns

The number of times packets were dropped because the bus had more packets to deliver than the MPM could handle. This is a receive overrun condition which can happen when a storm occurs or when the switch is first powered up and many unknown MAC frames are being forwarded to the MPM.

MPM Total Memory (bytes)

The amount of total memory installed on the MPM.

MPM Free Memory (bytes)

The amount of free, or unused, memory available in the MPM.

MPM CPU Utilization (5 Seconds)

The amount of time, by percent, the MPM processor actually worked during the last 5 seconds.

MPM CPU Utilization (60 Seconds)

The amount of time, by percent, that the MPM processor actually worked during the last minute.

Power Supply 1 State

Valid states are OK, Not Present, and Bad. A power supply that has been turned off will be in the Bad state. If not installed, it will be in the Not Present state.

Power Supply 2 State

Valid states are OK, Not Present, and Bad. A power supply that has been turned off will be in the Bad state. If not installed, it will be in the Not Present state.

Temperature

Indicates whether the MPM temperature sensor detects overheating.

Temperature

The current temperature of the chassis, as reported by the MPM module, both in degrees Celsius and Fahrenheit.

Environment Tab (XOS Devices)

The Environment tab reports the status of the chassis power supplies and provides information on chassis temperature and flash memory. Each field is described below.

The Environment Tab

Device: NMS-test-148 (10.255.11.148)

General Modules System State **Environment** Physical Port Virtual Port Interfaces Statistics

Power Supply Status

Power Module 1: bad

Backup Power Supply: okay

Temperature

Sensor 0 (Heat) Sensor 1 (Cold)

MPM (Primary): okay okay

MPM (Secondary): okay okay

Chassis Temperature: 40.0/104.0 Deg. (C/F)

Flash Memory Update Mode

Write Cache to Flash Memory

Write to Cache Only

Apply Update Help

Click **Update** to poll the switch and refresh the screen with current information.

Click **Apply** to write changes to the switch. All changes take effect immediately.

Power Supply Status

The status of each power supply (Power Module 1 and Backup Power Supply) in the chassis can be reported as:

OK. The power supply is installed and functioning.

Not Present. A power supply is not installed.

Bad. The power supply has failed. (**Note:** If a power supply is turned off, it might be reported as bad.)

Unknown. Power supply not recognized.

Temperature

MPM (Primary) / MPM (Secondary). All MPM modules have a temperature sensor (Sensor 0) that detects temperatures over 50° C. In addition, MPM 1Gs and MPM2s also have a second temperature sensor (Sensor 1) that detects temperatures under 0° C. The temperature range of an MPM module can be reported as:

OK. The MPM is operating within the allowed temperature tolerance for heat or cold (under 50° C or over 0° C, respectively).

Too hot or Too cold. The MPM is operating outside the allowed temperature tolerance for heat or cold, respectively, and may fail.

Not Present. An MPM is not installed in the slot. Some switch models do not use an MPM; in this case, the MPM (Slot 2) field will always display Not Present.

Chassis Temperature. The current temperature of the chassis as reported by the primary MPM module, in degrees Celsius and Fahrenheit. (**Note:** Display of the chassis temperature is supported by selected hardware only.)

Flash Memory Update Mode

Caution: Before using this feature, be sure to read the information below. While this feature does give you flexibility about when and how configuration information is saved, it can also inadvertently lead to loss of configuration changes.

To save configuration changes you may select from the following options:

Write Cache to Flash Memory. This is the default setting. When switch configuration changes are made to any program within OmniVista, those changes are written to the switch's cache, then saved to the switch's flash memory. This prevents configuration changes from being lost during a reboot. However, if numerous configuration changes are being made, Write Cache to Flash Memory may not be the best option to select, in that it involves increased processing time as one configuration change after another is written to the switch. When a series of configuration changes is being made, the Write to Cache Only option may be preferred.

Note: While cache is being written to flash memory, the SNMP agent will not be able to communicate with OmniVista for approximately 30 seconds.

Write to Cache Only. Writes switch configuration changes only to the switch's temporary cache. This option allows you to omit the step of writing changes to flash memory. Write to Cache Only allows all programs within OmniVista to respond to SNMP sets and gets faster, with no timeouts due to compaction, thereby enhancing switch performance. To activate the Write to Cache Only option, select its radio button, then click Apply. This will force switch configuration changes to be written only to cache. However, in the event of a reboot anytime after the Apply button has been clicked, configuration changes will be lost.

Physical Port Tab (XOS Devices)

The Physical Port tab provides information on all physical ports on the switch. This information is retrieved from the MIB (Management Information Base) phyPortTable. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Physical Port Tab

| Slot | Port | Media Type | Description | Admin. Status | Oper. Status |
|------|------|------------|------------------------|---------------|--------------|
| 3 | 1 | Ethernet | ETHR Physical Port 254 | enable | portDown |
| 3 | 2 | Ethernet | ETHR Physical Port 253 | enable | portDown |
| 3 | 3 | Ethernet | ETHR Physical Port 252 | enable | portDown |
| 3 | 4 | Ethernet | ETHR Physical Port 251 | enable | portDown |
| 3 | 5 | Ethernet | ETHR Physical Port 250 | enable | portDown |
| 3 | 6 | Ethernet | ETHR Physical Port 249 | enable | portDown |
| 3 | 7 | Ethernet | ETHR Physical Port 248 | enable | portDown |
| 3 | 8 | Ethernet | ETHR Physical Port 247 | enable | portDown |
| 3 | 9 | Ethernet | ETHR Physical Port 246 | enable | portDown |
| 3 | 10 | Ethernet | ETHR Physical Port 245 | enable | portDown |
| 3 | 11 | Ethernet | ETHR Physical Port 244 | enable | portDown |
| 3 | 12 | Ethernet | ETHR Physical Port 243 | enable | portDown |
| 3 | 13 | Ethernet | ETHR Physical Port 242 | enable | portDown |
| 3 | 14 | Ethernet | ETHR Physical Port 241 | enable | portDown |
| 3 | 15 | Ethernet | ETHR Physical Port 240 | enable | portUp |
| 3 | 16 | Ethernet | ETHR Physical Port 239 | enable | portDown |
| 3 | 17 | Ethernet | ETHR Physical Port 238 | enable | portDown |
| 3 | 18 | Ethernet | ETHR Physical Port 237 | enable | portDown |

Click Update to poll the switch and refresh the screen with current information.

Slot/Port

The slot and port for which status is displayed.

MediaType

The physical type of the port.

Description

A description of the port.

Admin Status

The Administrative (Admin) status of the port: Enabled or Disabled. When the Admin status of a port is enabled, the port can receive and transmit data as long as a cable is connected and no physical or operational problems exist. When the Administrative Status of a port is disabled, 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 Administrative Status is enabled.

OperStatus

The operational status of the port: PortUp, PortDown, or Unknown.

Virtual Port Tab (XOS Devices)

The Virtual Port tab, shown below, displays information and status for all virtual ports on an XOS switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Virtual Port Tab

Device: NMS-194-1 (10.255.11.194)

General Modules Environment Physical Port **Virtual Port** Interfaces Statistics

Virtual Port Status Table 44/44

| Slot | Port | Service | Instance | Number | Group | Admin. State | Oper. Status | Description |
|------|------|---------|----------|--------|-------|--------------|--------------|--------------------|
| 3 | 1 | bridge | 1 | 13 | 1 | enable | portDown | Virtual port (#13) |
| 3 | 2 | bridge | 1 | 14 | 1 | enable | portDown | Virtual port (#14) |
| 3 | 3 | bridge | 1 | 15 | 1 | enable | portDown | Virtual port (#15) |
| 3 | 4 | bridge | 1 | 16 | 1 | enable | portDown | Virtual port (#16) |
| 3 | 5 | bridge | 1 | 17 | 1 | enable | portDown | Virtual port (#17) |
| 3 | 6 | bridge | 1 | 18 | 1 | enable | portDown | Virtual port (#18) |
| 3 | 7 | bridge | 1 | 19 | 1 | enable | portDown | Virtual port (#19) |
| 3 | 8 | bridge | 1 | 20 | 1 | enable | portDown | Virtual port (#20) |
| 3 | 9 | bridge | 1 | 21 | 1 | enable | portDown | Virtual port (#21) |
| 3 | 10 | bridge | 1 | 22 | 1 | enable | portDown | Virtual port (#22) |
| 3 | 11 | bridge | 1 | 23 | 1 | enable | portDown | Virtual port (#23) |
| 3 | 12 | bridge | 1 | 24 | 1 | enable | portDown | Virtual port (#24) |
| 3 | 13 | bridge | 1 | 37 | 1 | enable | portDown | Virtual port (#37) |
| 3 | 14 | bridge | 1 | 38 | 1 | enable | portDown | Virtual port (#38) |
| 3 | 15 | bridge | 1 | 39 | 1 | enable | portDown | Virtual port (#39) |
| 3 | 16 | bridge | 1 | 40 | 1 | enable | portDown | Virtual port (#40) |

Update Help

Click Update to poll the switch and refresh the screen with current information.

Slot/Port

The slot and port on which the virtual port resides.

Service

The service provided by this virtual port (Router, Bridge, Trunk, AtmTrunk, AtmLANE, 1483 Scaling, etc.).

Instance

The specific instance of this Slot/Port/Service. For most interface types the instance will always be 1. ATM-connected ports are an exception.

Number

A unique number that identifies this virtual port instance within the switch.

Group

The group to which this port belongs. Group 1 is the default group.

Admin State

The administrative status of this port: Enable or Disable.

Oper Status

The operational status of this port: portUp or portDown.

Description

An alphanumeric string that describes the instance of this port.

MAC Address

The MAC address of this port.

Encapsulation

The kind of frames that are sent out this port. If translation is necessary, incoming frames are translated to the format displayed here before being sent out this port. This field can display as:

Switch. Translations are governed by the vportSwitchTable.

Media Default. Translations are governed by the vportSwitchDefaultTable, which is indexed by media type of port.

EthII-LLC. Valid only for Ethernet/Ethernet LANE ports, this translates all IPX encapsulations except 802.2 to Ethertype.

LLC. Valid for all media types, this translates all IPX encapsulations to 802.2 LLC.

Snap-LLC. Valid for all media types, this translates all IPX encapsulations except 802.2 LLC to SNAP.

EthII. Valid only for Ethernet/Ethernet LANE ports, this translates all IPX encapsulations including 802.2 to Ethertype.

Snap. Valid for all media types, this translates all IPX encapsulations including 802.2 LLC to SNAP.

Bridge Protocol

The type of Bridge Protocol supported. For Ethernet ports, the default Bridge Protocol is Transparent. Some non-Ethernet ports (such as Token Ring) can also have a Bridge Protocol of SourceRouting or SRTransparent (Source Routing Transparent).

Bridge Mode

The Bridge Mode can display as AutoSwitch, ForceBridge, or ForceSwitch, as explained below:

AutoSwitch. The switch automatically switches the port between Optimized Device Switching mode and Spanning Tree Bridge mode depending on the number of MAC addressees seen attached to the port. Initially the port is placed in Optimized Device Switching mode, but once the switch detects more than one MAC address attached to the port, it switches the port into Spanning Tree Bridge mode.

ForceBridge. The port acts as a standard Spanning Tree 802.1d bridge port. It forwards Spanning Tree BPDU frames out the port. When frames are received, Spanning Tree BPDUs are processed, and Spanning Tree dynamically controls the forwarding state. If flooding occurs, all frames destined for unknown MAC addresses, broadcast addresses, and multicast addresses are sent to all ports on the same VLAN.

ForceSwitch. This mode is appropriate when only one MAC address, such as a file or mail server, is attached to the port. Since only one device is attached, no Spanning Tree BPDUs are sent and the port is always in the forwarding state. Unknown unicast frames are not flooded. However, if the port is set to ForceSwitch and more than one MAC address or Spanning Tree BPDU is detected, the port is automatically changed to a Spanning Tree Bridge port and an SNMP trap is generated to that effect.

Manual Mode

This field displays the port's manual Spanning Tree status. Manual Spanning Tree configuration is primarily designed to allow override of forwarding or blocking on ports regardless of their Spanning Tree state (IBM Spanning Tree active or IEEE 802.1d Spanning Tree active). However, manual configuration of Spanning Tree can also be used to enable IEEE 802.1d Spanning Tree on virtual ports that do not support IBM Spanning Tree (such as Ethernet and FDDI ports). When such a port is present in a group that has been assigned the IBM Spanning Tree algorithm, the switch automatically overrides the assignment and does not run any Spanning Tree on the port. You can enable IEEE 802.1d Spanning Tree on such a port by setting this field to Dynamic. This field can display:

Dynamic. Manual mode is disabled. Spanning Tree -- either IBM Spanning Tree or IEEE 802.1d Spanning Tree, as appropriate -- is enabled. If this port is an Ethernet port in a Group that is assigned IBM Spanning Tree, IEEE 802.d Spanning Tree will be enabled for this port.

OverrideFwd. Do not allow forwarding at this port.

OverrideBlock. Do not allow blocking at this port.

Switch Timer

When the Bridge Mode field (described above) displays AutoSwitch, the value in the Switch Timer field defines the timeout period, in seconds, before a port operating in Spanning Tree Bridge Mode converts to Optimized Device Switching Mode. When set to AutoSwitch, the port initially operates in Optimized Device Switching Mode but switches to Spanning Tree Bridge Mode if more than one MAC address is detected. The port will switch back to AutoSwitch mode after the timeout period displayed here. The default value for this field is 60 seconds. When this field is set to zero (0), immediate switching between the two modes occurs.

Flood Limit

The flood limit enables the "tuning" of a virtual port to limit the flooding of broadcast, multicast, and unknown destination packets. This feature is useful for controlling broadcast storms on the network. While each network is different, in general the amount of flooded traffic represents a relatively small percentage of network traffic.

The flood limit is actually a "transmit credit" that is issued every five seconds. When a packet is flooded on the port, the size of the packet, in bytes, is decremented from the current credit value. The credit value is the value displayed in this field multiplied by five. An additional credit, of the value displayed in this field multiplied by five, is allocated to the virtual port every five seconds. If the credit value falls below zero, all flooded packets are discarded until another credit is allocated. Flood limit checking is disabled if a flood limit value of zero (0) displays. The flood limit default value is 192,000 bytes per second, which equates to a transmit credit of 960,000 bytes every five seconds.

Interfaces Tab (XOS Devices)

The Interfaces tab provides status for all interfaces on the switch. This information is retrieved from the MIB (Management Information Base) ifTable. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Interfaces Tab

| Index | Description | Type | MTU | Speed | Physical Address |
|-------|--|--------------------|------|---------|-------------------|
| 1 | vn254 | PROP VIRTUAL | 1500 | 960 Mbs | 00 d0 95 4e 32 26 |
| 2 | pcn0 | ETHERNET-CSMA/CD | 1500 | 10 Mbs | 00 d0 95 4e 32 20 |
| 3 | lo0 | SOFTWARE-LOOP-BACK | 4096 | 10 Mbs | |
| 3001 | 3/1 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 20 |
| 3002 | 3/2 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 21 |
| 3003 | 3/3 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 22 |
| 3004 | 3/4 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 23 |
| 3005 | 3/5 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 24 |
| 3006 | 3/6 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 25 |
| 3007 | 3/7 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 26 |
| 3008 | 3/8 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 27 |
| 3009 | 3/9 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 28 |
| 3010 | 3/10 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 29 |
| 3011 | 3/11 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 2a |
| 3012 | 3/12 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 2b |
| 3013 | 3/13 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 2c |
| 3014 | 3/14 10/100Mb Ethernet CSMA/CD interface | ETHERNET-CSMA/CD | 1518 | 100 Mbs | 00 d0 95 5f a9 2d |

Click Update to poll the switch and refresh the screen with current information.

Index

A unique value that identifies this interface.

Description

A textual description of the interface.

Type

The type of the interface, identified according to the physical or link protocol(s) immediately "below" the network layer in the protocol stack.

MTU

The size, in octets, of the largest datagram that can be sent or received on this interface. This is the size of the largest network datagram that can be transmitted on interfaces used for transmitting network datagrams.

Speed

An estimate of the interface's current bandwidth. Speed is displayed in bits-per-second if less than 1,000,000 bits-per-second. Speeds of 1,000,000 bits-per-second or greater are displayed in terms of Mbs (megabits-per-second). If an interface does not vary in bandwidth, or if no accurate estimation can be made, the nominal bandwidth is displayed in this field.

Physical Address

The interface address at the protocol layer (the layer immediately "below" the network layer). This field displays no value for interfaces that do not have such an address (for example, a serial line).

Admin State

The administrative state of the interface: Up indicates the interface is administratively enabled to pass packets; Down indicates the interface is administratively disabled from passing packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Oper Status

The operational state of the interface: Up indicates the interface is able to pass packets; Down indicates the interface is not able to pass packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Last Change

The amount of time since the interface entered its current operational state. This field will display a zero if the current operational state was entered prior to the last reinitialization of the application.

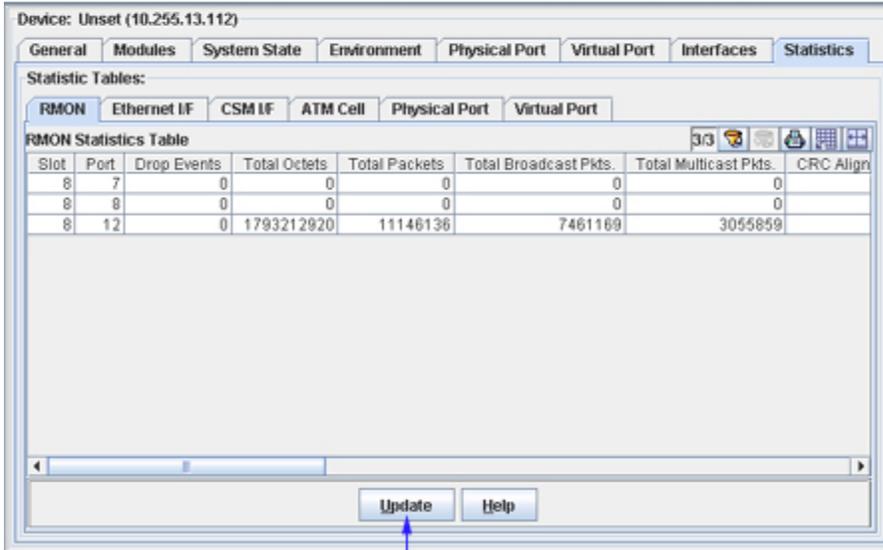
Out Queue

The length of the output packet queue (in packets).

RMON Statistics (XOS Devices)

The RMON Tab displays remote monitoring statistics for all Ethernet ports on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The RMON (Remote Monitoring) Statistics Tab



Click **Update** to poll the switch and refresh the screen with current information.

Slot/Port

The slot and port for which statistics are displayed.

Drop Events

The total number of events during which packets were dropped by the probe due to lack of resources. Note that this number is not necessarily the number of packets dropped; it is the number of times this condition has been detected.

Total Octets

The total number of octets of data received, including those in bad packets (excluding framing bits but including FCS -- Frame Check Sequence -- octets). This value can be used as a reasonable estimate of Ethernet utilization. If greater precision is desired, the **Total Packets** and **Total Octets** fields should be sampled before and after a common interval. In the equation below, the differences in the sampled values are Packets and Octets, respectively, and the number of seconds in the interval is Interval. These values are used to calculate the Ethernet utilization as follows:

$$\text{Utilization} = \frac{\text{Packets} * (8.6 + 6.4) + (\text{Octets} * .8)}{\text{Interval} * 10,000}$$

The result of this equation is the value Utilization which is the percent utilization of the Ethernet segment on a scale of 0 to 100 percent (per RFC 1757).

Total Packets

The total number of packets (including bad packets, broadcast packets, and multicast packets) received.

Total Broadcast Pkts

The total number of good packets received that were directed to the broadcast address. Note that this value does not include multicast packets.

Total Multicast Pkts

The total number of good packets received that were directed to a multicast address. Note that this value does not include packets directed to the broadcast address.

CRC Align Errors

The total number of packets received that had a length between 64 and 1518 octets, inclusive (excluding framing bits but including FCS octets), but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).

Undersized Pkts

The total number of packets received that were less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.

Oversized Pkts

The total number of packets received that were more than 1518 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.

Fragments

The total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Note that it is entirely normal for the value in this field to increment. This is because the **Fragments** field counts both runts (which are normal occurrences due to collisions) and noise hits.

Jabbers

The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).

Note that this definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.

Collisions

The best estimate of the total number of collisions on this Ethernet segment. The value returned depends on the location of the RMON probe. Section 8.2.1.3 (10BASE-5) and section 10.3.1.3 (10BASE-2) of IEEE standard 802.3 states that a station must detect a collision, in the receive mode, if three or more stations are transmitting simultaneously. A repeater port must detect a collision when two or more stations are transmitting simultaneously. Thus a probe placed on a repeater port could record more collisions than would a probe connected to a station on the same segment.

Probe location plays a much smaller role when considering 10BASE-T. 14.2.1.4 (10BASE-T) of IEEE standard 802.3 defines a collision as the simultaneous presence of signals on the DO and RD circuits (transmitting and receiving at the same time). A 10BASE-T station can only detect collisions when it is transmitting. Thus probes placed on a station and a repeater should report the same number of collisions.

Note also that an RMON probe inside a repeater should ideally report collisions between the repeater and one or more other hosts (transmit collisions as defined by IEEE 802.3k) plus receiver collisions observed on any coax segments to which the repeater is connected.

Pkts 64 Octets

The total number of packets received (including bad packets) that were 64 octets in length (excluding framing bits but including FCS octets).

Pkts 65-127 Octets

The total number of packets received (including bad packets) that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).

Pkts 128-255 Octets

The total number of packets received (including bad packets) that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).

Pkts 256-511 Octets

The total number of packets received (including bad packets) that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).

Pkts 512-1023 Octets

The total number of packets received (including bad packets) that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).

Pkts 1024-1518 Octets

The total number of packets received (including bad packets) that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).

Ethernet Interface Statistics (XOS Devices)

The Ethernet I/F tab lists statistics for each Ethernet interface in the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below. Note that discontinuities can occur in statistics values upon reinitialization of the system.

The Ethernet Interfaces Tab

Device: Unset (10.255.13.112)

General Modules System State Environment Physical Port Virtual Port Interfaces Statistics

Statistic Tables:

RMON Ethernet IF CSM IF ATM Cell Physical Port Virtual Port

Ethernet Interface Statistics Table

| Slot | Port | Index | Type | Rx Octets | Tx Octets | Rx Unicast Pkts | Tx Unicast Pkts | Rx Nc |
|------|------|-------|------------------|-----------|-----------|-----------------|-----------------|-------|
| 5 | 1 | 5001 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 5 | 2 | 5002 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 1 | 7001 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 2 | 7002 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 3 | 7003 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 4 | 7004 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 5 | 7005 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 6 | 7006 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 7 | 7007 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 8 | 7008 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 9 | 7009 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 10 | 7010 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 11 | 7011 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |
| 7 | 12 | 7012 | ETHERNET-CSMA/CD | 0 | 0 | 0 | 0 | |

Update Help

Click Update to poll the switch and refresh the screen with current information.

Slot/Port

The slot and port for which statistics are displayed.

Index

A unique value that identifies this interface.

Type. The type of the interface, identified according to the physical or link protocol(s) immediately "below" the network layer in the protocol stack.

Rx Octets

The number of octets, or bytes, received on this interface.

Tx Octets

The number of octets, or bytes, transmitted from this interface.

Rx Unicast Pkts

The number of unicast packets received on this interface.

Tx Unicast Pkts

The number of unicast packets transmitted from this interface.

Rx Non-Unicast Pkts

The number of non-unicast packets received on this interface.

Tx Non-Unicast Pkts

The number of non-unicast packets transmitted from this interface.

Rx I/F Discards

The number of frames received on this interface discarded due to lack of buffer space.

Tx I/F Discards

The number of frames that could not be transmitted from this interface due to lack of buffer space.

Rx I/F Errors

The number of frames received on this interface discarded due to errors.

Tx I/F Errors

The number of frames that could not be transmitted from this interface due to errors.

Unknowns

The number of frames received on this interface with an unknown protocol.

CSM Interface Statistics (XOS Devices)

The CSM I/F Tab displays CSM interface statistics for a physical CSM port. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The CSM Interfaces Tab

Device: Unset (10.255.13.112)

General Modules System State Environment Physical Port Virtual Port Interfaces Statistics

Statistic Tables:

RMON Ethernet IF CSM IF ATM Cell Physical Port Virtual Port

CSM Interface Statistics Table 11/11

| Slot | Port | Rx Cells | Tx Cells | Rx Cells CLP=0 | Rx Cells CLP=1 | Marked EFCI Cells | Marked GCRA |
|------|------|----------|----------|----------------|----------------|-------------------|-------------|
| 2 | 1 | 7427647 | 7427578 | 7427647 | 0 | 0 | 0 |
| 4 | 1 | 0 | 0 | 0 | 0 | 0 | 0 |
| 4 | 2 | 0 | 0 | 0 | 0 | 0 | 0 |
| 4 | 3 | 0 | 0 | 0 | 0 | 0 | 0 |
| 4 | 4 | 0 | 0 | 0 | 0 | 0 | 0 |
| 4 | 5 | 0 | 0 | 0 | 0 | 0 | 0 |
| 4 | 6 | 0 | 0 | 0 | 0 | 0 | 0 |
| 4 | 7 | 0 | 0 | 0 | 0 | 0 | 0 |
| 4 | 8 | 0 | 0 | 0 | 0 | 0 | 0 |
| 9 | 1 | 0 | 0 | 0 | 0 | 0 | 0 |
| 9 | 2 | 0 | 0 | 0 | 0 | 0 | 0 |

Update Help

Click Update to poll the switch and refresh the screen with current information.

Slot/Port

The slot number of the CSM module and the port number for which statistics are displayed. Each row in the table gives information for a single CSM port.

Rx Cells

The total number of cells received on this port since the last initialization of the switch. This count includes all received cells (data, management, and discarded).

Tx Cells

The total number of cells transmitted from this CSM port since the last initialization of the switch. This count includes all transmitted cells.

Rx Cells CLP=0

The number of ATM cells received on this port with the CLP bit set to 0. Cells with the CLP bit set to 0 (CLP=0) are high priority and cells with a CLP bit set to 1 (CLP=1) are low priority. Refer to ATM Traffic Management, above, for further information.

Rx Cells CLP=1

The number of ATM cells received on this port with the CLP bit set to 1. Cells with the CLP bit set to 0 (CLP=0) are high priority and cells with a CLP bit set to 1 (CLP=1) are low priority. Refer to ATM Traffic Management, above, for further information. Because of the switch's policing algorithms, there is a higher probability of CLP1 cells being discarded than CLP0 cells.

Marked EFCI Cells

The number of ATM cells in which the Explicit Forward Congestion Indication (EFCI) bit is set. The EFCI notification is used in conjunction with backward RM cells so that the destination can notify the source that there is congestion on the path to the destination.

Marked GCRA Cells

The number of ATM cells marked by the policing GCRA for violating the traffic contract for CLP=0+1 cells.

Total Discard Cells

The total number of cells discarded at this interface due to congestion, policing, and cells with unknown VPIs or VCIs.

Dx Congestion CLP=0

The number of CLP0 (high priority) cells discarded at this interface due to congestion.

Dx Congestion CLP=1

The number of CLP1 (low priority) cells discarded at this interface due to congestion.

Dx GCRA(A) CLP=0

The total number of CLP0 (high priority) cells discarded at this interface due to policing on CLP=0+1 cells by the first GCRA, or leaky bucket.

Dx GCRA(A) CLP=1

The total number of CLP1 (low priority) cells discarded at this interface due to policing on CLP=0+1 cells by the first GCRA, or leaky bucket.

Dx GCRA(B) CLP=0

The total number of CLP0 (high priority) cells discarded at this interface due to policing on CLP0 cells by the second GCRA, or leaky bucket.

Dx GCRA(B) CLP=1

The total number of CLP1 (low priority) cells discarded at this interface due to policing on CLP=0+1 cells by the second GCRA, or leaky bucket.

Unknown VP/VC Cells

The number of cells received on this interface with a VPI/VCI (Virtual Path Identifier/Virtual Channel Identifier) combination that does not correspond to the VPI/VCI combination of any virtual circuit on this physical interface.

Unknown VPI

The last unknown VPI (Virtual Path Identifier) received; that is, the last VPI received on this interface that does not correspond to the VPI of any virtual circuit on this interface. Please note that this parameter is not currently supported.

Unknown VCI

The last unknown VCI (Virtual Channel Identifier) received; that is, the last VCI received on this interface that does not correspond to the VCI of any virtual circuit on this interface. Please note that this parameter is not currently supported.

UniType

The type of UNI (User-to-Network Interface) used on this interface. This field may display the following:

Public. Public User-to-Network Interface. This interface is used for connections to public ATM service carrier switches, such as those used by Telcos.

Private. Private User-to-Network Interface. This interface is used for private UNI uplinks. Such a port would connect either directly to an ATM workstation, LAN switch, or ATM attached router.

PNNI. This interface supports PNNI (Private Network-to-Network Interface) version 1.0 ATM routing, which includes support for a single peer group mapping. PNNI is a dynamic routing protocol that is capable of establishing switched virtual connections based on ATM End System requests. PNNI is also capable of managing connections that use preconfigured static routes. Static routes are used by the Interim Inter-Switch Signaling Protocol (IISP), which is an ATM static routing protocol.

IISP-Net. This interface supports an IISP (Interim Interswitch Signaling Protocol) network connection. Typically an IISP interface would be part of an intermediate ATM node that did not support the PNNI routing protocol, and would be used primarily for establishing static routes using the IISP protocol. An IISP interface must be configured to be either the user side or the network side. This is important because only one side of a link can be the network side, which allocates all the Virtual Circuits.

IISP-USER. This interface supports an IISP (Interim Interswitch Signaling Protocol) user-side connection. Please refer to **IISP-Net**, above, for further information on IISP.

UniVersion

The version of the UNI (User-to-Network Interface) used on this interface. The switch is compliant with ATM Forum UNI specifications versions 3.0 and 3.1. This field may display the following:

UNI 30. This interface is compliant with ATM Forum UNI 3.0.

UNI 31. This interface is compliant with ATM Forum UNI 3.1.

UniIISP. This interface is compliant with IISP signaling. IISP can imitate UNI 3.0 or UNI 3.1 signaling. Please refer to **IISP-Net**, above, for further information on IISP.

Rx Remaining Bandwidth

This field displays the remaining bandwidth available on this interface through which connections can be created and cells can be received.

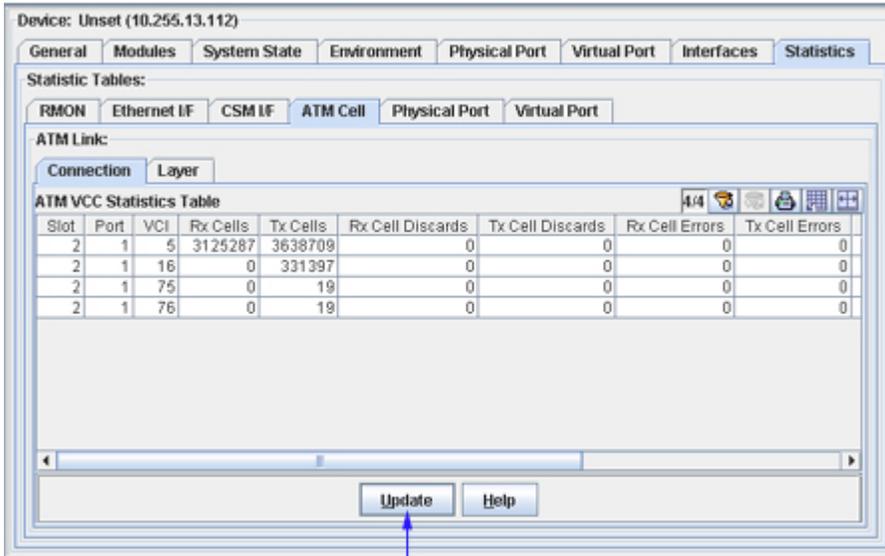
Tx Remaining Bandwidth

This field displays the remaining bandwidth available on this interface through which connections can be created and cells can be transmitted.

ATM Cell Statistics (XOS Devices)

The ATM Cell Tab displays ATM cell statistics for all ports on ASM submodules. You can view either ATM **Connection** statistics or ATM **Layer** statistics. Note that the same fields display on the Connection tab and the Layer tab. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The ATM Cell Tab



Click **Update** to poll the switch and refresh the screen with current information.

Connection or Layer Statistics

Slot

A unique value which identifies this hsm board slot.

Port

A unique value which identifies this atm submodule.

VCI (Connection statistics only)

A unique identifier associated with the virtual channel.

Rx Cells

The total number of cells within a SDU (service data unit) that were successfully received.

Tx Cells

The total number of cells within a SDU (service data unit) that were successfully transmitted.

Rx Cell Discards

The total number of receive cells discarded due to SDU discards. When an SDU is discarded, the cells that compose the SDU are counted and this statistic is incremented accordingly.

Tx Cell Discards

The total number of transmit cells discarded due to SDU discards. When an SDU is discarded, the cells that compose the SDU are counted and this statistic is incremented accordingly.

Rx Cell Errors

The total number of cells within receive SDUs that had one or more of the following errors: invalid format, frame larger than the Rx buffer, frame larger than the maximum size allowed on this virtual connection, invalid size, or CRC errors. For each SDU with errors, the number of cells within that SDU are counted and this statistic is incremented accordingly.

Tx Cell Errors

The total number of cells within transmit SDUs that had one or more of the following errors: invalid format, frame larger than the Rx buffer, frame larger than the maximum size allowed on this virtual connection, invalid size, or CRC errors. For each SDU with errors, the number of cells within that SDU are counted and this statistic is incremented accordingly.

Rx Cell No Buffers

The total number of receive cells that were discarded due to insufficient space in the frame buffer. Note that the cells counted in this statistic are not included in the Discard statistic.

Tx Cell No Buffers

The total number of transmit cells that were discarded due to insufficient space in the frame buffer. Note that the cells counted in this statistic are not included in the Discards or Errors statistic.

Rx Cell Trash

The number of cells that never left the ATM physical layer. These cells were discarded by the SAR buffer due to a lack of reassembly buffer space. Note that the cells counted in this statistic are not included in the Discards or Errors statistic.

Physical Port Statistics (XOS Devices)

The Physical Port Tab displays statistics for all physical ports on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Physical Port Tab

Device: Unset (10.255.13.112)

Statistic Tables:

Physical Port Statistics Table

| Slot | Port | Rx Frames | Tx Frames | Rx Octets | Tx Octets | Rx Unicast Pkts | Tx Unicast Pkts | Rx Non-Uni |
|------|------|-----------|-----------|------------|-----------|-----------------|-----------------|------------|
| 8 | 2 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 8 | 3 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 8 | 4 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 8 | 5 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 8 | 6 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 8 | 7 | 2 | 11 | 66 | 996 | 0 | 6 | 6 |
| 8 | 8 | 0 | 11 | 0 | 996 | 0 | 6 | 6 |
| 8 | 9 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 8 | 10 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 8 | 11 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 8 | 12 | 11700490 | 257582 | 1933236... | 48381101 | 666499 | 256630 | |
| 9 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

Update Help

Click **Update** to poll the switch and refresh the screen with current information.

Slot/Port

The slot and port for which statistics are displayed.

Rx Frames

The number of frames received on this port.

Tx Frames

The number of frames transmitted from this port.

Rx Octets

The number of octets, or bytes, received on this port.

Tx Octets

The number of octets, or bytes, transmitted from this port.

Rx Unicast Pkts

The number of unicast packets received on this port.

Tx Unicast Pkts

The number of unicast packets transmitted from this port.

Rx Non-Unicast Pkts

The number of non-unicast packets received on this port.

Tx Non-Unicast Pkts

The number of non-unicast packets transmitted from this port.

Rx Buffer Discards

The number of frames received on this port discarded due to lack of buffer space.

Tx Buffer Discards

The number of frames that could not be transmitted from this port due to lack of buffer space.

Rx Error Discards

The number of frames received on this port discarded due to errors.

Tx Error Discards

The number of frames that could not be transmitted from this port due to errors.

Virtual Port Statistics (XOS Devices)

The Virtual Port Tab displays statistics for all virtual ports on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Virtual Port Tab

Device: Unset (10.255.13.112)

Statistic Tables:

Virtual Port Statistics Table

| Slot | Port | Service | Instance | Number | Group | Rx Frames | Tx Frames | Rx Octets | Tx Octets | Rx Uni |
|------|------|---------|----------|--------|-------|-----------|-----------|-----------|-----------|--------|
| 2 | 1 | bridge | 1 | 48 | 5 | 0 | 8 | 0 | 0 | 0 |
| 2 | 1 | bridge | 2 | 49 | 6 | 0 | 8 | 0 | 0 | 0 |
| 5 | 1 | bridge | 1 | 2 | 1 | 0 | 0 | 0 | 0 | 0 |
| 5 | 2 | bridge | 1 | 3 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 1 | bridge | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 2 | bridge | 1 | 4 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 3 | bridge | 1 | 6 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 4 | bridge | 1 | 7 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 5 | bridge | 1 | 5 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 6 | bridge | 1 | 9 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 7 | bridge | 1 | 10 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 8 | bridge | 1 | 11 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 9 | bridge | 1 | 12 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 10 | bridge | 1 | 13 | 1 | 0 | 0 | 0 | 0 | 0 |

Update Help

Click Update to poll the switch and refresh the screen with current information.

Slot/Port

The physical slot and port numbers for this virtual port instance.

Service

The function of this virtual port: Router, Bridge, Trunk, AtmTrunk, AtmLANE, etc.

Instance

The specific instance of this Slot/Port/Service. For most interface types the instance will always be 1. ATM-connected ports are an exception.

Number

A unique number that identifies this virtual port instance within the physical switch.

Group

The Group to which this port belongs. Group 1 is the default group.

Rx Frames

The total number of frames received on this port since the last time the switch was initialized.

Tx Frames

The total number of frames transmitted from this port since the last time the switch was initialized.

Rx Octets

The total number of Octets, or bytes, received on this port since the last time the switch was initialized. This statistic includes the data and Frame Relay header fields, but does not include CRC or flag characters.

Tx Octets

The total number of Octets, or bytes, sent on this port since the last time the switch was initialized. This statistic includes the data and Frame Relay header fields, but does not include CRC or flag characters.

Rx Unicast Pkts

The total number of subnetwork unicast packets received from this port.

Tx Unicast Pkts

The total number of subnetwork unicast packets transmitted from this port.

Rx Non-Unicast Pkts

The total number of non-unicast packets received from this port.

Tx Non-Unicast Pkts

The total number of non-unicast packets transmitted from this port.

Rx Buffer Discards

The number of inbound frames discarded from this port due to overruns of the receive queue.

Tx Buffer Discards

The number of outbound frames discarded from this port due to overruns of the transmit queue.

Rx Error Discards

The number of inbound frames discarded from this port due to errors.

Tx Error Discards

The number of outbound frames discarded from this port due to errors.

Flood Limit Discards

The number of outbound frames discarded from this port due to the flood limit being exceeded.

General Tab (6024 Devices)

The General tab for 6024 devices provides general system information, as explained below. To change a configurable parameter, edit the field as desired and then click **Apply** to write the change to the switch. All changes take effect immediately.



The General Tab

Click **Update** to poll the switch and refresh the screen with current information.

Click **Apply** to write changes to the switch. All changes take effect immediately.

Name

A user-defined name for this switch.

Description

A description of the switch as defined by the manufacturer.

Location

A user-defined description of the switch's physical location.

Contact

A user-defined parameter stating who is responsible for this switch.

Up Time

The period of time that has elapsed since the switch was last rebooted.

Interfaces Tab (6024 Devices)

The Interfaces tab provides status for all interfaces on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Interfaces Tab

Click Update to poll the switch and refresh the screen with current information.

| Index | Description | Type | MTU | Speed | Physical Address | Admin. State | Oper. Status |
|-------|-------------------|------------------|------|---------|-------------------|--------------|--------------|
| 1 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 21 | up | up |
| 2 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 22 | up | down |
| 3 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 10 Mbs | 00 94 5a 58 29 23 | up | down |
| 4 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 10 Mbs | 00 94 5a 58 29 24 | up | down |
| 5 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 25 | up | up |
| 6 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 10 Mbs | 00 94 5a 58 29 26 | up | down |
| 7 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 27 | up | up |
| 8 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 28 | up | up |
| 9 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 29 | up | down |
| 10 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 2a | up | up |
| 11 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 2b | up | down |
| 12 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 2c | up | up |
| 13 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 10 Mbs | 00 94 5a 58 29 2d | up | down |
| 14 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 10 Mbs | 00 94 5a 58 29 2e | up | down |
| 15 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 10 Mbs | 00 94 5a 58 29 2f | up | down |
| 16 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 10 Mbs | 00 94 5a 58 29 30 | up | down |
| 17 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 31 | up | down |
| 18 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 94 5a 58 29 32 | up | up |
| 19 | Ethernet 10/100MB | ETHERNET-CSMA/CD | 1500 | 10 Mbs | 00 94 5a 58 29 33 | up | down |

Index

A unique value that identifies this interface.

Description

A textual description of the interface.

Type

The type of the interface, identified according to the physical or link protocol(s) immediately "below" the network layer in the protocol stack.

MTU

The size, in octets, of the largest datagram that can be sent or received on this interface. This is the size of the largest network datagram that can be transmitted on interfaces used for transmitting network datagrams.

Speed

An estimate of the interface's current bandwidth. Speed is displayed in bits-per-second if less than 1,000,000 bits-per-second. Speeds of 1,000,000 bits-per-second or greater are displayed in terms of Mbs (megabits-per-second). If an interface does not vary in bandwidth, or if no accurate estimation can be made, the nominal bandwidth is displayed in this field.

Physical Address

The interface address at the protocol layer (the layer immediately "below" the network layer). This field displays no value for interfaces that do not have such an address (for example, a serial line).

Admin State

The administrative state of the interface: Up indicates the interface is administratively enabled to pass packets; Down indicates the interface is administratively disabled from passing packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Oper Status

The operational state of the interface: Up indicates the interface is able to pass packets; Down indicates the interface is not able to pass packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Last Change

The amount of time since the interface entered its current operational state. This field will display a zero if the current operational state was entered prior to the last re-initialization of the application.

Out Queue

The length of the output packet queue (in packets).

Ports Tab (6024 Devices)

The Ports tab devices provides information on the physical ports on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Ports Tab

| Unit Id | Port Id | Port Type | Admin Speed and Mode | Oper Speed and Mode | Admin Flow Control | Oper Flow Control |
|---------|---------|---------------|----------------------|---------------------|--------------------|-------------------|
| 1 | 1 | hundredBaseTX | autoNegotiation | fullDuplex100 | disabled | none |
| 1 | 2 | hundredBaseTX | autoNegotiation | fullDuplex100 | disabled | none |
| 1 | 3 | hundredBaseTX | autoNegotiation | halfDuplex10 | disabled | none |
| 1 | 4 | hundredBaseTX | autoNegotiation | halfDuplex10 | disabled | none |
| 1 | 5 | hundredBaseTX | autoNegotiation | fullDuplex100 | disabled | none |
| 1 | 6 | hundredBaseTX | autoNegotiation | halfDuplex10 | disabled | none |
| 1 | 7 | hundredBaseTX | autoNegotiation | fullDuplex100 | disabled | none |
| 1 | 8 | hundredBaseTX | autoNegotiation | halfDuplex100 | disabled | none |
| 1 | 9 | hundredBaseTX | autoNegotiation | fullDuplex100 | disabled | none |
| 1 | 10 | hundredBaseTX | autoNegotiation | fullDuplex100 | disabled | none |
| 1 | 11 | hundredBaseTX | autoNegotiation | fullDuplex100 | disabled | none |
| 1 | 12 | hundredBaseTX | autoNegotiation | fullDuplex100 | disabled | none |
| 1 | 13 | hundredBaseTX | autoNegotiation | halfDuplex10 | disabled | none |
| 1 | 14 | hundredBaseTX | autoNegotiation | halfDuplex10 | disabled | none |
| 1 | 15 | hundredBaseTX | autoNegotiation | halfDuplex10 | disabled | none |
| 1 | 16 | hundredBaseTX | autoNegotiation | halfDuplex10 | disabled | none |

Click **Update** to poll the switch and refresh the screen with current information.

Port ID

An ID number that identifies the port within this switch.

Port Type

The type of the port.

Admin Speed and Mode

The speed and duplex mode to which the port is set administratively. The value in this field may be **halfDuplex1000** (1000 Mbps and half duplex mode), **fullDuplex1000** (1000 Mbps and full duplex mode), or **autoNegotiation** (allow the switch to negotiate duplex mode and speed with the other end of connection).

Oper Speed and Mode

The speed and duplex mode at which the port is actually operating. The value in this field may be **halfDuplex1000** or **fullDuplex1000**.

Admin Flow Control

The administrative state of flow control for the port: either **enabled** or **disabled**. When flow control is enabled, and the port is operating in halfDuplex mode, the backPressure flow control mechanism is used. When flow control is enabled, and the port is operating in fullDuplex mode, the IEEE 802.3x flow control mechanism is used. Flow control can eliminate frame loss by "blocking" traffic from end stations or segments connected directly to the switch when switch buffers fill.

Oper Flow Control

The type of flow control the port is actually using during operation. This field may display the following values:

backPressure. The backPressure flow control mechanism is in use. The backPressure flow control mechanism is used when flow control is administratively enabled and the port is operating in halfDuplex mode at 1000 Mbps.

dot3xFlowControl. The IEEE 802.3x flow control mechanism is in use. The IEEE 802.3x flow control mechanism is used when flow control is administratively enabled and the port is operating in fullDuplex mode at 1000 Mbps.

none. Flow control is disabled.

General Tab (6100 Devices)

The General tab for 6100 devices provides general system information and general chassis information, as explained below. To change any configurable parameter, edit the field as desired and then click **Apply** to write the change to the switch. All changes take effect immediately.

The General Tab

Device: test (10.255.11.224)

General Ports Interfaces

System

Name: Test Description: Alcatel OmniStack 6124

Location: Lab Contact: John x-4679

Up Time: 15 days 4 hours 5 minutes 47 seconds AMAP: Active

Chassis Information Table

| Unit Id | Hardware Version | Firmware Version | Agent Hardware Version | Agent Firmware Version |
|---------|------------------|------------------|------------------------|------------------------|
| 1 | V4.0 | V1.29 | V2.0 (850 CPU) | V3.40.31 |

Apply Update Help

Click Update to poll the switch and refresh the screen with current information.

Click Apply to write changes to the switch. All changes take effect immediately.

System Parameters

Name

A user-defined name for this switch.

Description

A factory-defined description of the switch.

Location

A user-defined description of the switch's physical location.

Contact

A user-defined statement identifying the person or organization responsible for the switch.

Up Time

The period of time that has elapsed since the switch was last rebooted.

AMAP

Set this field to **Active** or **Inactive** to enable or disable the AMAP protocol on this switch. By default, AMAP is enabled. AMAP is a proprietary protocol that learns the connections and links between switches in the list of All Discovered Devices. This information is used to create a graphical display of network links when a network region or subnet is viewed. If you disable AMAP, this switch's connections and links will not be displayed.

Chassis Information Parameters

Unit ID

An ID number that identifies the switch.

Hardware Version

The hardware version of the main board.

Firmware Version

The version of the firmware on the main board.

Agent Hardware Version

The hardware version of the agent board.

Agent Firmware Version

The version of the firmware on the agent board.

POST Code Version

The version of the POST (Power On Self Test) code on the agent board.

Port Count

The total number of ports on the switch, including expansion slots.

Power Status

Indicates whether the switch is using internalPower, redundantPower, or both internalAndRedundantPower.

Expansion Slot 1

The type of module installed in Expansion Slot 1. If no module is installed, **notPresent** displays.

Expansion Slot 2

The type of module installed in Expansion Slot 2. If no module is installed, **notPresent** displays.

Role in System

Indicates whether the switch is functioning as the **master**, **backupMaster**, or **slave**.

Interfaces Tab (6100 Devices)

The Interfaces tab provides status for all interfaces on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Interfaces Tab

| Index | Description | Type | MTU | Speed | Physical Address | Admin. |
|-------|-------------------|-------------------|------|---------|-------------------|--------|
| 1 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a a1 | up |
| 2 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a a2 | up |
| 3 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a a3 | up |
| 4 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a a4 | up |
| 5 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 100 Mbs | 00 30 f1 15 1a a5 | up |
| 6 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a a6 | up |
| 7 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a a7 | up |
| 8 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a a8 | up |
| 9 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 100 Mbs | 00 30 f1 15 1a a9 | up |
| 10 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a aa | up |
| 11 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a ab | up |
| 12 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a ac | up |
| 13 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a ad | up |
| 14 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 100 Mbs | 00 30 f1 15 1a ae | up |
| 15 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 100 Mbs | 00 30 f1 15 1a af | up |
| 16 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a b0 | up |
| 17 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a b1 | up |
| 18 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a b2 | up |
| 19 | Ethernet 10/100MB | ETHERNET-C SMA/CD | 1500 | 10 Mbs | 00 30 f1 15 1a b3 | up |

Click Update to poll the switch and refresh the screen with current information.

Index

A unique value that identifies this interface.

Description

A textual description of the interface.

Type

The type of the interface, identified according to the physical or link protocol(s) immediately "below" the network layer in the protocol stack.

MTU

The size, in octets, of the largest datagram that can be sent or received on this interface. This is the size of the largest network datagram that can be transmitted on interfaces used for transmitting network datagrams.

Speed

An estimate of the interface's current bandwidth. Speed is displayed in bits-per-second if less than 1,000,000 bits-per-second. Speeds of 1,000,000 bits-per-second or greater are displayed in terms of Mbs (megabits-per-second). If an interface does not vary in bandwidth, or if no accurate estimation can be made, the nominal bandwidth is displayed in this field.

Physical Address

The interface address at the protocol layer (the layer immediately "below" the network layer). This field displays no value for interfaces that do not have such an address (for example, a serial line).

Admin State

The administrative state of the interface: Up indicates the interface is administratively enabled to pass packets; Down indicates the interface is administratively disabled from passing packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Oper Status

The operational state of the interface: Up indicates the interface is able to pass packets; Down indicates the interface is not able to pass packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Last Change

The amount of time since the interface entered its current operational state. This field will display a zero if the current operational state was entered prior to the last re-initialization of the application.

Out Queue

The length of the output packet queue (in packets).

Ports Tab (6100 Devices)

The Ports tab provides information on the physical ports on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Ports Tab

| Unit Id | Port Id | Port Type | Admin Speed and Mode | Oper Speed and Mode | Admin Flow |
|---------|---------|---------------|----------------------|---------------------|------------|
| 1 | 1 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 2 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 3 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 4 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 5 | hundredBaseTX | autoNegotiation | fullDuplex100 | enabled |
| 1 | 6 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 7 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 8 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 9 | hundredBaseTX | autoNegotiation | fullDuplex100 | enabled |
| 1 | 10 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 11 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 12 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 13 | hundredBaseTX | fullDuplex100 | halfDuplex10 | enabled |
| 1 | 14 | hundredBaseTX | fullDuplex100 | fullDuplex100 | enabled |
| 1 | 15 | hundredBaseTX | autoNegotiation | fullDuplex100 | enabled |
| 1 | 16 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 17 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 18 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |
| 1 | 19 | hundredBaseTX | autoNegotiation | halfDuplex10 | enabled |

Click Update to poll the switch and refresh the screen with current information.

Port ID

An ID number that identifies the port within this switch.

Port Type

The type of the port.

Admin Speed and Mode

The speed and duplex mode to which the port is set administratively. The value in this field may be **halfDuplex1000** (1000 Mbps and half duplex mode), **fullDuplex1000** (1000 Mbps and full duplex mode), or **autoNegotiation** (allow the switch to negotiate duplex mode and speed with the other end of connection).

Oper Speed and Mode

The speed and duplex mode at which the port is actually operating. The value in this field may be **halfDuplex1000** or **fullDuplex1000**.

Admin Flow Control

The administrative state of flow control for the port: either **enabled** or **disabled**. When flow control is enabled, and the port is operating in halfDuplex mode, the backPressure flow control mechanism is used. When flow control is enabled, and the port is operating in fullDuplex mode, the IEEE 802.3x flow control mechanism is used. Flow control can eliminate frame loss by "blocking" traffic from end stations or segments connected directly to the switch when switch buffers fill.

Oper Flow Control

The type of flow control the port is actually using during operation. This field may display the following values:

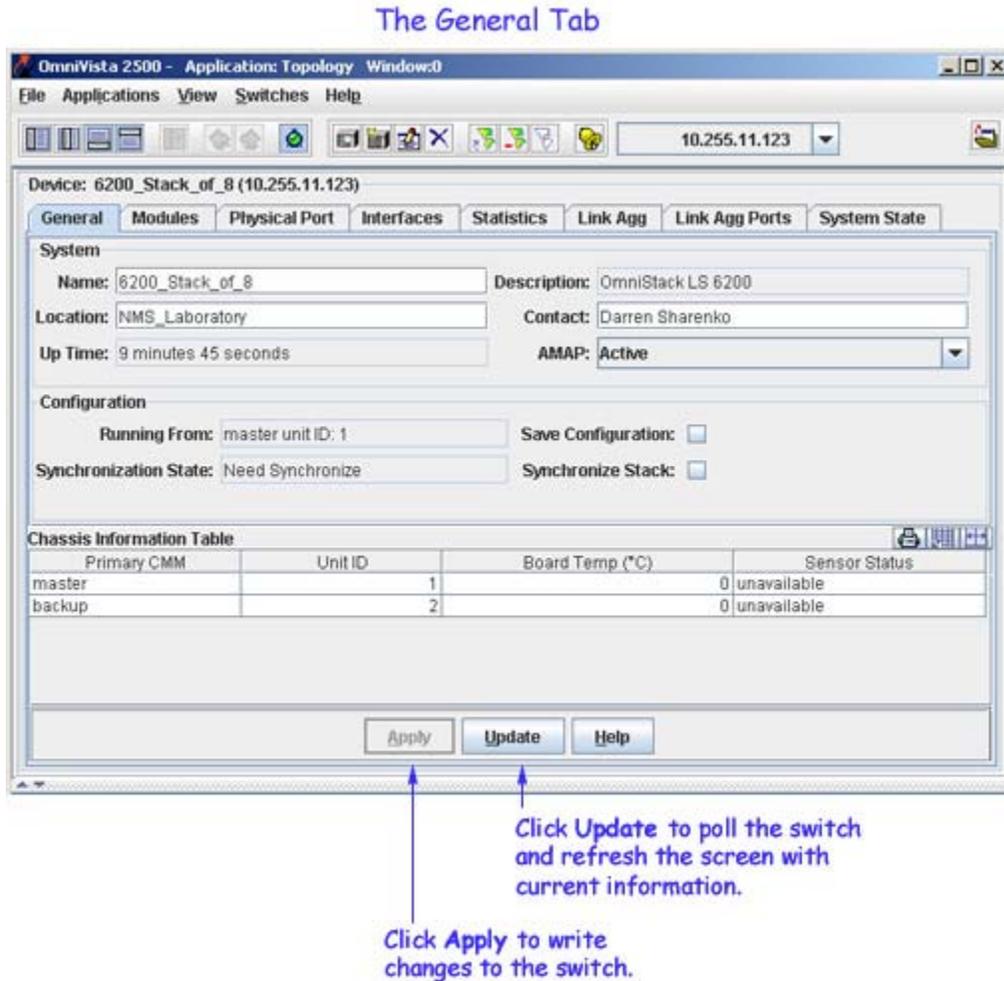
backPressure. The backPressure flow control mechanism is in use. The backPressure flow control mechanism is used when flow control is administratively enabled and the port is operating in halfDuplex mode at 1000 Mbps.

dot3xFlowControl. The IEEE 802.3x flow control mechanism is in use. The IEEE 802.3x flow control mechanism is used when flow control is administratively enabled and the port is operating in fullDuplex mode at 1000 Mbps.

none. Flow control is disabled.

General Tab (6200 Devices)

The General tab for 6200 devices displays general system information and specific chassis information. It also enables you to start and stop the AMAP protocol and to save and synchronize switch configuration files. You can change any configurable parameter by editing the field desired and then clicking **Apply** to write the change to the switch. All changes take effect immediately.



System Parameters

Name

A user-defined name for this switch.

Description

A factory-defined description of the switch's software.

Location

A user-defined description of the switch's physical location.

Contact

A user-defined statement identifying the person or organization responsible for the switch.

Up Time

The period of time that has elapsed since the switch was last rebooted.

AMAP

Set this field to **Active** or **Inactive** to enable or disable the AMAP protocol on this switch. By default, AMAP is enabled. AMAP is a proprietary protocol that learns the connections and links between switches in the list of All Discovered Devices. This information is used to create a graphical display of network links when a network region or subnet is viewed.

Configuration Parameters

The fields in the section are used to save the configuration files, and synchronize the switches in a stack (stacked configurations only) .



Running From

Displays the unit in the stack which is acting as the Master.

Save Configuration

Issues a "Save Configuration" command to the device to save the configuration files.

Synchronization State (Stacked Configuration Only)

Displays the synchronization state for a stacked configuration (Synchronized/Need Synchronize). This field is only visible in stacked configurations.

Synchronize Stack (Stacked Configuration Only)

Issues a "Synchronize Stack" command to synchronize the image files (and boot files, if necessary) of all of the devices in a stack. This checkbox is activated if the Synchronization State is "Need Synchronize" (a new switch was added to the stack, new image files were added to the master). This command copies the latest image files to the working directory on each switch in the stack. To activate the new image files you must re-boot the Master Switch from the working directory by right-clicking on the switch in the device tree and selecting **Reboot>From Working**. This field is only visible in stacked configurations.

Note: This command does not synchronize configuration files between the master and backup devices. Configuration files are automatically synchronized each time a new command is issued by the user. The master unit synchronizes both the running (RAM) and startup configurations (static).

Chassis Information Parameters

| Primary CMM | Unit ID | Board Temp.(°C) | Sensor Status |
|-------------|---------|-----------------|---------------|
| master | 1 | 0 | unavailable |
| backup | 2 | 0 | unavailable |

Primary CMM

This field identifies the switch that is currently functioning as the primary CMM. In a stacked configuration, the primary switch is identified as "Master", the secondary switch is identified as "Backup".

Unit ID

The role of the switch in the stack. "1" identifies the Master switch, "2" identifies the Backup switch (if applicable).

Board Temp (Degrees Celsius)

The current reading of the board temperature sensor, in degrees Celsius, for this chassis.

Sensor Status

Sensor status of the Master and Backup (if applicable) switches in the stack.

Note: Not all fields display for all devices. If a field is not applicable to a device it is not displayed.

Modules Tab (6200 Devices)

The Modules tab lists the hardware modules installed in the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each column is described below.

The Modules Tab

| Slot | Name | Type | Description | HwRevision | SerialNumber | FwVersion | SwRevision |
|------|--------------------|--------|-------------------------|------------|--------------|-----------|------------|
| 1 | OmniStack LS 6224U | backup | Alcatel 24F+4G fiber | 00.00.02 | G2050013 | 1.0.0.11 | 1.0.2.36 |
| 2 | OmniStack LS 6224U | master | Alcatel 24F+4G fiber | 00.00.02 | G2050012 | 1.0.0.11 | 1.0.2.38 |
| 3 | OmniStack LS 6212P | slave | Alcatel 12F+4G with POE | 00.00.04 | G3154972 | 1.0.0.11 | 1.0.2.38 |
| 4 | OmniStack LS 6212P | slave | Alcatel 12F+4G with POE | 00.00.04 | G3154967 | 1.0.0.11 | 1.0.2.38 |
| 5 | OmniStack LS 6224U | slave | Alcatel 24F+4G fiber | 00.00.02 | G2050024 | 1.0.0.11 | 1.0.2.38 |
| 6 | OmniStack LS 6224 | slave | Alcatel 24F+4G Non PoE | 00.00.01 | F2950354 | 1.0.0.11 | 1.0.2.38 |
| 7 | OmniStack LS 6224U | slave | Alcatel 24F+4G fiber | 00.00.02 | G2050018 | 1.0.0.11 | 1.0.2.38 |
| 8 | OmniStack LS 6224P | slave | Alcatel 24F+4G with PoE | 00.00.01 | F3952018 | 1.0.0.11 | 1.0.2.38 |

Slot

The slot in which the module is installed.

Name

The name of the module

Type

The factory-defined physical type of the module.

Description

A description of the module.

Hw Revision

The current revision level of the module hardware

Serial Number

Serial number of the module.

Fw Version

The module's firmware version. All modules should use the same firmware version.

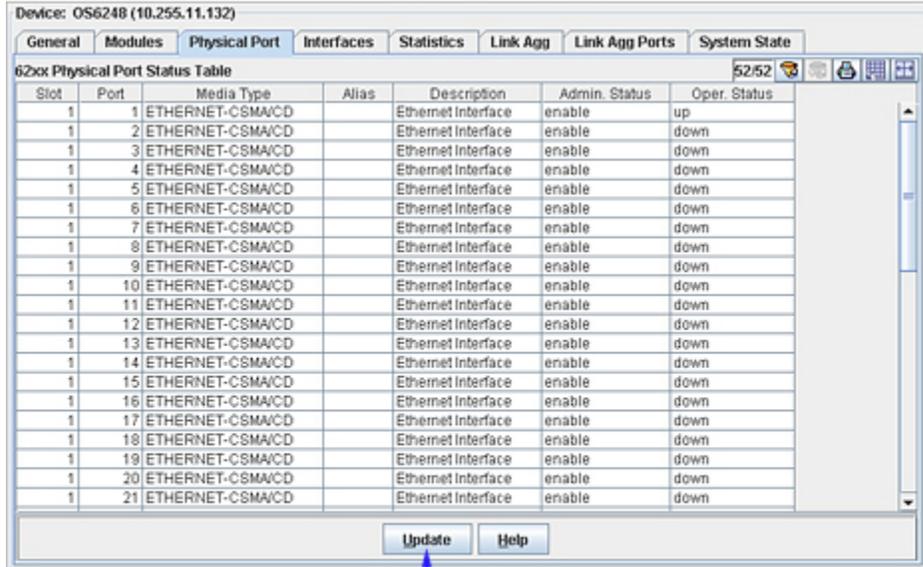
Sw Revision

The module's software version. All modules should use the same software version.

Physical Port Tab (6200Devices)

The Physical Port tab provides information on all physical ports on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Physical Port Tab



Device: OS6248 (10.255.11.132)

General Modules **Physical Port** Interfaces Statistics Link Agg Link Agg Ports System State

620x Physical Port Status Table 52/52

| Slot | Port | Media Type | Alias | Description | Admin. Status | Oper. Status |
|------|------|-------------------|-------|--------------------|---------------|--------------|
| 1 | 1 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | up |
| 1 | 2 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 3 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 4 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 5 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 6 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 7 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 8 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 9 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 10 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 11 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 12 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 13 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 14 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 15 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 16 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 17 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 18 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 19 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 20 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |
| 1 | 21 | ETHERNET-C SMA/CD | | Ethernet Interface | enable | down |

Update Help

Click Update to poll the switch and refresh the screen with current information.

Slot/Port

The slot and port for which status is displayed.

Media Type

The physical type of the port.

Alias

The user-defined alias for the port.

Description

A description of the port.

Admin Status

The Administrative (Admin) status of the port: **up** or **down**. When the Admin status of a port is enabled, the port can receive and transmit data as long as a cable is connected and no physical or operational problems exist. When the Administrative Status of a port is disabled, 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 Administrative Status is enabled.

OperStatus

The operational status of the port: **portUp**, **portDown**, or **unknown**.

The Interfaces Tab (6200 Devices)

The Interfaces tab provides information about all active interfaces on the OS6200 switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

Note: In stacked configuration, the table lists all possible physical ports and this leads to slow response in reading the table.

The Interfaces Tab

| Index | Alias | Description | Type | MTU | Speed | Physical Address | Admin. State | Oper. Status |
|-------|-------|--------------------|------------------|------|---------|-------------------|--------------|--------------|
| 1 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c a1 | up | up |
| 2 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c a2 | up | down |
| 3 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c a3 | up | down |
| 4 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c a4 | up | down |
| 5 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c a5 | up | down |
| 6 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c a6 | up | down |
| 7 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c a7 | up | down |
| 8 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c a8 | up | down |
| 9 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c a9 | up | down |
| 10 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c aa | up | down |
| 11 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c ab | up | down |
| 12 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c ac | up | down |
| 13 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c ad | up | down |
| 14 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c ae | up | down |
| 15 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c af | up | down |
| 16 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c b0 | up | down |
| 17 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c b1 | up | down |
| 18 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c b2 | up | down |
| 19 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c b3 | up | down |
| 20 | | Ethernet Interface | ETHERNET-CSMA/CD | 1500 | 100 Mbs | 00 12 cf 0b 8c b4 | up | down |

Click Update to poll the switch and refresh the screen with current information.

Index

A unique value that identifies this interface internally.

Description

A description of the interface that usually includes the name of the manufacturer, the name of the product, and the version of the interface's hardware/software.

Type

A description of the type of the interface.

MTU

The size, in octets, of the largest packet that can be sent or received on the interface.

Speed

An estimate of the interface's current bandwidth. Speed is displayed in bits-per-second if less than 1,000,000 bits-per-second. Speeds of 1,000,000 bits-per-second or greater are displayed in terms of Mbs

(megabits-per-second). If an interface does not vary in bandwidth, or if no accurate estimation can be made, the nominal bandwidth is displayed in this field.

Physical Address

The physical address of the interface at its protocol sublayer. For 802.x interfaces, the physical address is a MAC address. No physical address displays for interfaces in loopback mode nor for serial interfaces.

Admin. State

The administrative state of the interface: **up**, **down**, or **testing**. Admin state **up** indicates the interface is administratively enabled to pass packets; **down** indicates the interface is administratively disabled from passing packets; **testing** indicates the interface is in a test mode and cannot pass operational packets. All interfaces are initialized with the admin state **down**. After initialization, either in response to explicit management action or stored configuration data, the admin state of an interface to changed to **up** or **testing** (or may remain **down**).

Oper. Status

The current operational status of the interface: **up**, **down**, **testing**, **unknown**, **dormant**, **notPresent**, or **lowerLayerDown**.

up. The interface is ready to transmit and receive packets.

down. The interface is either administratively disabled or there is a fault that prevents it from going to the **up** state.

testing. The interface is in a test mode and cannot pass operational packets.

dormant. The interface is waiting for external actions (such as a serial line waiting for an incoming connection).

notPresent. The interface has missing components (typically hardware components).

lowerLayerDown. The interface is down due to the state of lower-layer interfaces.

If an interface's administrative state is **down** its operational status will also be **down**. When the administrative state 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 going **up**; or, the operational status will remain

Last Change

The value of sysUpTime when the interfaces table (ifTable) was last changed because a new entry was created or an existing entry was deleted. (The sysUpTime MIB variable reports the time period that has elapsed since the switch was last initialized.) If the interfaces table was not changed since the last re-initialization of OmniVista, no value will display in this field.

Out Queue

The length of the packet output queue, in packets.

RMON Statistics (6200 Devices)

The RMON Statistics tab, displays RMON (Remote Monitoring) statistics information for 6200 devices. The interfaces that are not connected will be filtered out from the view. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

Note: In a stacked configuration, the table contains all possible physical ports and this leads to slow response reading the table.

The RMON Statistics Tab

| Slot | Port | Drop Events | Total Octets | Total Packets | Total Broadcast Pkts. | Total Multicast Pkts. | CRC Align Errors |
|------|------|-------------|--------------|---------------|-----------------------|-----------------------|------------------|
| 1 | 1 | 0 | 2572079455 | 63620978 | 28492690 | 1785321 | 3 |
| 1 | 2 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 3 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 4 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 5 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 6 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 7 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 8 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 9 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 10 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 11 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 12 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 13 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 14 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 15 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 16 | 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 17 | 0 | 0 | 0 | 0 | 0 | 0 |

Click Update to poll the switch and refresh the screen with current information.

Slot and Port

The slot and port for which RMON statistics are displayed.

Drop Events

The total number of occasions that packets were dropped by the probe due to lack of resources. Note that the value in this field is not necessarily the number of packets dropped; it is the number of times this condition was detected.

Total Octets

The total number of octets received, including those in bad packets. The count includes FCS (frame check sequence) octets but excludes framing bits. The value in this field can be used as a reasonable estimate of 10 megabit Ethernet utilization. If greater precision is desired, the **Total Octets** and **Total Packets** values should be sampled before and after a common interval. In the following equation, the differences in the sampled values are *Octets* and *Pkts*, respectively, and the number of seconds in the common interval is *Interval*. The result of this equation is the value *Utilization* which is the percent utilization of the Ethernet segment on a scale of 0 to 100 percent.

$$\text{Utilization} = \frac{\text{Pkts} * (9.6 + 6.4) + (\text{Octets} * .8)}{\text{Interval} * 10,000}$$

Total Packets

The total number of packets received, including bad packets, broadcast packets, and multicast packets.

Total Broadcast Pkts

The total number of good packets received that were directed to the broadcast address. Note that this value does not include multicast packets.

Total Multicast Pkts

The total number of good packets received that were directed to a multicast address. Note that this value does not include packets directed to the broadcast address.

CRC Align Errors

The total number of packets received with a length between 64 and 1518 octets, inclusive (excluding framing bits but including FCS [frame check sequence] octets), which had either of the following errors:

a bad frame check sequence with an integral number of octets, which is an FCS error, or

a bad frame check sequence with a non-integral number of octets, which is an alignment error.

Undersized Pkts

The total number of packets received that were less than 64 octets in length, excluding framing bits but including FCS (frame check sequence) octets, and were otherwise well formed.

Oversized Pkts

The total number of packets received that were longer than 1518 octets, excluding framing bits but including FCS (frame check sequence) octets, and were otherwise well formed.

Fragments

The total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS [frame check sequence] octets), which had either of the following errors:

a bad frame check sequence with an integral number of octets, which is an FCS error, or

a bad frame check sequence with a non-integral number of octets, which is an alignment error.

Note that it is entirely normal for the count in this field to increment, because it includes both runt packets (which are a normal occurrence due to collisions) and noise hits.

Jabbers

The total number of packets received that were longer than 1518 octets (excluding framing bits but including FCS [frame check sequence] octets), which had either of the following errors:

a bad frame check sequence with an integral number of octets, which is an FCS error, or

a bad frame check sequence with a non-integral number of octets, which is an alignment error.

Note that this definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.

Collisions

The best estimate of the total number of collisions on this Ethernet segment. The value returned depends on the location of the RMON probe. Section 8.2.1.3 (10BASE-5) and section 10.3.1.3 (10BASE-2) of IEEE standard 802.3 states that a station, when in receive mode, must detect a collision if three or more stations are transmitting simultaneously. A repeater port must detect a collision when two or more stations are transmitting simultaneously. Thus, a probe placed on a repeater port could record more collisions than would a probe connected to a station on the same segment.

Probe location plays a much smaller role when considering 10BASE-T. Section 14.2.1.4 (10BASE-T) of IEEE standard 802.3 defines a collision as the simultaneous presence of signals on the DO and RD circuits (transmitting and receiving at the same time). A 10BASE-T station can only detect collisions when it is transmitting. Thus, a probe placed on a station and a probe placed on a repeater should report the same number of collisions.

Note that an RMON probe inside a repeater should ideally report collisions between the repeater and one or more other hosts (per the IEEE 802.3k definition of transmit collisions) plus receiver collisions observed on any coax segments to which the repeater is connected.

Pkts 64 Octets

The total number of packets received, including bad packets, that were 64 octets in length. The count includes FCS (frame check sequence) octets but excludes framing bits.

Pkts 65-127 Octets

The total number of packets received, including bad packets, that were between 65 and 127 octets in length, inclusive. The count includes FCS (frame check sequence) octets but excludes framing bits.

Pkts 128-255 Octets

The total number of packets received, including bad packets, that were between 128 and 255 octets in length, inclusive. The count includes FCS (frame check sequence) octets but excludes framing bits.

Pkts 256-511 Octets

The total number of packets received, including bad packets, that were between 256 and 511 octets in length, inclusive. The count includes FCS (frame check sequence) octets but excludes framing bits.

Pkts 512-1023 Octets

The total number of packets received, including bad packets, that were between 512 and 1023 octets in length, inclusive. The count includes FCS (frame check sequence) octets but excludes framing bits.

Pkts 1024-1518 Octets

The total number of packets received, including bad packets, that were between 1024 and 1518 octets in length, inclusive. The count includes FCS (frame check sequence) octets but excludes framing bits.

Ethernet Statistics (6200 Devices)

The Ethernet I/F tab displays the Ethernet statistics for the OS6200 device. The interfaces that are not connected will be filtered out from the view. Only the physical interfaces are displayed. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below. Note that discontinuities can occur in statistics values upon reinitialization of the system.

Note: In a stacked configuration, the table lists all the possible physical ports and this leads to slow response in reading the table.

The Ethernet I/F Statistics Tab

| Slot | Port | Index | Type | Rx Octets | Tx Octets | Rx Unicast Pkts | Tx Unicast Pkts | Rx IF Discards |
|------|------|-------|-------------------|------------|-----------|-----------------|-----------------|----------------|
| 1 | 1 | 1 | ETHERNET-CSDMA/CD | 2572104690 | 742198290 | 33342782 | 5522126 | 0 |
| 1 | 2 | 2 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 3 | 3 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 4 | 4 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 5 | 5 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 6 | 6 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 7 | 7 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 8 | 8 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 9 | 9 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 10 | 10 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 11 | 11 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 12 | 12 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 13 | 13 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 14 | 14 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 15 | 15 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 16 | 16 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |
| 1 | 17 | 17 | ETHERNET-CSDMA/CD | 0 | 0 | 0 | 0 | 0 |

Click Update to poll the switch and refresh the screen with current information.

Slot and Port

The slot and port of the interface.

Index

A unique value that identifies the interface internally.

Type

The type of the interface.

Rx Octets

The total number of octets received on the interface, including framing characters.

Tx Octets

The total number of octets transmitted out of the interface, including framing characters.

Rx Unicast Pkts

The total number of unicast packets received on this interface and delivered to a higher layer. This value does not include packets addressed to a multicast or broadcast address.

Tx Unicast Pkts

The total number of unicast packets that higher-level protocols requested be transmitted from this interface, including packets that were discarded or not sent. This value does not include packets addressed to a multicast or broadcast address at this sublayer.

Rx I/F Discards

The number of received packets that were discarded even though no errors were detected in the packets that would have prevented them from being delivered to a higher-layer protocol. One possible reason for discarding such packets would be the need to free buffer space.

Tx I/F Discards

The number of outbound packets that were discarded even though no errors were detected in the packets that would have prevented them from being transmitted. One possible reason for discarding such packets would be the need to free buffer space.

Rx I/F Errors

The number of received packets that contained errors preventing them from being delivered to a higher-layer protocol.

Tx I/F Errors

The number of outbound packets that could not be transmitted because of errors.

Unknowns

The number of received packets that were discarded because of an unknown or unsupported protocol.

Link Agg Tab (6200 Devices)

The Link Agg tab displays all active Link Aggregate information for the OS6200 device. Link Aggregate interfaces that are not connected are from the view.

Link aggregation is a way of combining multiple physical links between two switches into one logical link. The aggregate group operates within Spanning Tree as one virtual port and can provide more bandwidth than a single link. It also provides redundancy. If one physical link in the aggregate group goes down, link integrity is maintained.

There are two types of aggregate groups: static and dynamic. Static aggregate groups are manually configured on the switch with static links. Dynamic groups are set up on the switch but they aggregate links as necessary according to the Link Aggregation Control Protocol (LACP). OmniVista's Link Agg tab provides information about each link aggregation group defined on the switch. Each field in the tab is described below.



Click Update to poll the switch and refresh the screen with current information.

Number

A reference number assigned when the link aggregation group was created.

Size

The maximum number of links that may belong to this link aggregation group.

Name

The name of the link aggregation group. This is an alphanumeric string up to 255 characters long.

Description

The standard MIB name for this link aggregate group.

LACP Type

The type of this link aggregation group. **lacpOff** means the group is static. **lacpOn** means the group is dynamic and is using the LACP protocol. (LACP is the Link Aggregation Control Protocol.)

Admin State

The administrative state of this link aggregation group: either **enable** (the group is active and is able to aggregate links) or **disable** (the group is inactive). The group's administrative state is configured by the network administrator.

Oper State

The current operational state of this link aggregation group: either **up** (the group is operational) or **down** (the group is not operational). This field may also display **logicPortCreatFailed** or **qReservationFailed**.

Selected Ports

The number of ports that could possibly attach to this link aggregation group at the moment.

Attached Ports

The number of ports actually attached to this link aggregation group at the moment.

Primary Port

The slot/port number of the primary port in the link aggregation group used to send BPDUs and flooding frames. The switch uses the first port to join the group as the primary port. If the first port to join the group is no longer part of the group, the switch automatically assigns another port in the group to be the primary port.

MAC Address

The MAC address assigned to this link aggregation group.

Actor System ID

The MAC address for the local port associated with a dynamic link aggregation group, which is used as a unique identifier for the system that contains this link aggregation group.

Actor System Priority

A value from 0 - 65535 that indicates the priority value associated with the Actor System ID. This defines the priority of the switch's dynamic aggregate group in relation to other aggregate groups

Actor Admin Key

The administrative key value configured for the dynamic aggregate group. Possible values are 0 - 65535.

Actor Oper Key

The current operational value of the key for the dynamic link aggregation group.

Partner System ID

The MAC address of the remote aggregate group to which this aggregate group is attached. A value of zero indicates that there is no known partner. If the group is manually configured, the value in this field is assigned by the local system.

Partner System Priority

The priority of the remote system to which the aggregation group is attached. Possible values are 0 - 65535. If the group is manually configured, the value in this field is assigned by the local system.

Partner Admin Key

The administrative key for the aggregation group's remote partner. Possible values are 0 - 65535. If the group is manually configured, the value in this field is assigned by the local system. The administrative key may differ from the operational key.

Partner Oper key

The operational key of the remote system to which the aggregation group is attached. If the group is manually configured, the value in this field is assigned by the local system.

Link Agg Ports Tab (6200 Devices)

The Link Agg Ports tab displays all active Link Aggregate Ports information for the OS6200 device. Link Aggregate Ports that are not connected are filtered from the view. Each field is described below.

The Link Agg Ports Tab

| Slot/Port | Aggregate ID | Actor Priority | Actor ID | Actor Admin Key | Actor Oper Key | Partner Priority | Partner Oper |
|-----------|--------------|----------------|---------------|-----------------|----------------|------------------|--------------|
| 1/1 | LAG-1 | 1 | 0012cf.0b8ca0 | 1000 | 1000 | 0 | |
| 1/2 | LAG-1 | 1 | 0012cf.0b8ca0 | 1000 | 1000 | 0 | |
| 1/3 | LAG-1 | 1 | 0012cf.0b8ca0 | 1000 | 1000 | 0 | |
| 1/4 | LAG-1 | 1 | 0012cf.0b8ca0 | 1000 | 1000 | 0 | |

Click Update to poll the switch and refresh the screen with current information.

Slot/Port

The slot and port number of a port in the link aggregation group.

Aggregate ID

The ID of the static aggregate group to which the port is attached. This field does not apply to dynamic aggregate groups. The **Aggregate ID** can be any value from **-1** to **31**. The **-1** value displays when this field is not significant.

Admin State

The administrative state of this port: either **enable** (the port is ready to pass packets) or **disable** (the port is administratively disabled). The port's administrative state is configured by the network administrator.

Oper State

The operational status of the port: either **up** (the port is passing traffic), **down** (the port is unable to pass traffic) **notAttached** (the port is not attached to the aggregate group), or **notAggregable** (the port cannot be aggregated, perhaps because the key is not set or is incorrect).

Port State

The current aggregation status of the port. When a port is attached to a group, **attached** will display in this field. Other possible port states are **created**, **configurable**, **configured**, **selected**, and **reserved**.

Link State

The operational status of the link: **up** or **down**.

Primary

This field displays **yes** if the port is the primary port in the aggregate group and displays **no** if it is not. This field may also display **notSignificant**.

Actor System ID

The System ID (i.e., the MAC address) of the system that contains this port.

Actor System Priority

A value from **0 - 255** that defines the priority value associated with the Actor's System ID.

Actor Admin Key

The actor administrative key value for this port.

Actor Oper Key

The current operational value of the actor key.

Partner Admin System ID

The administrative MAC address associated with the remote partner's system ID. This value is used along with Partner Admin System Priority, Partner Admin Key, Partner Admin Port, and Partner Admin Port Priority to manually configure aggregation.

Partner Oper System Priority

The operational priority of the remote system to which this port is attached.

Partner Admin Key

The administrative value of the key for the remote partner. This value is used along with Partner Admin System Priority, Partner Admin System ID, Partner Admin Port, and Partner Admin Port Priority to manually configure aggregation.

Partner Oper Key

The current operational value of the key for the protocol partner.

Selected Agg ID

The Aggregator ID associated with the dynamic aggregate group to which the port is attached. Zero indicates that this port has not selected an aggregate group, either because it is in the process of detaching from a group or because there is no suitable group available for it to select.

Attach Agg ID

The Aggregator ID associated with the dynamic aggregate group to which the port is attached. Zero indicates that this port is not currently attached to a group.

Actor Port

The port number locally assigned to this port. The port number is communicated in Link Aggregation Control Protocol Data Units (LACPDU) as the Actor_Port (a read-only value).

Actor Port Priority

The actor priority value assigned to the port. The actor priority value can range from 0 - 255.

Partner Admin Port

The administrative value of the port number for the protocol partner. This value is used along with Partner

Admin System Priority, Partner Admin System ID, Partner Admin Key, and Partner Admin Port Priority to manually configure aggregation.

Partner Oper Port

The operational port number assigned to the port by the port's protocol partner.

Partner Admin Port Priority

The administrative port priority of the protocol partner. This value is used along with Partner Admin System Priority, Partner Admin System ID, Partner Admin Key, and Partner Admin Port to manually configure aggregation.

Partner Oper Port Priority

The priority value assigned to this port by the partner.

Actor Admin State

The administrative state of the port. The Actor Admin State is a string of eight bits that correspond to the administrative values of Actor_State, as transmitted by the Actor in Link Aggregation Control Protocol Data Units (LACPDU)s. The bits of Actor Admin State are as follows:

The first bit corresponds to bit 0 of Actor_State, which is Activity. When this bit is set, the dynamic aggregate group is able to exchange LACPDU frames.

The second bit corresponds to bit 1 of Actor_State, which is Timeout. When this bit is set, a short timeout is used for LACPDU frames. When this bit is disabled, a long timeout is used for LACPDU frames.

The third bit corresponds to bit 2 of Actor_State, which is Aggregation. When this bit is set, the system considers this port to be a potential candidate for aggregation. If this bit is not enabled, the system considers the port to be individual (it can only operate as a single link).

The fourth bit corresponds to bit 3 of Actor_State, which is Synchronization. The system always determines the value of this bit. When bit 3 is set by the system, the port is allocated to the correct dynamic aggregation group. If this bit is not set by the system, the port is not allocated to the correct dynamic aggregation group.

The fifth bit corresponds to bit 4 of Actor_State, which is Collecting. The system always determines the value of this bit. When bit 4 is set by the system, incoming LACPDU frames are collected from the individual ports that make up the dynamic aggregate group.

The sixth bit corresponds to bit 5 of Actor_State, which is Distributing. The system always determines the value of this bit. When bit 5 is set by the system, distributing outgoing frames on the port is disabled.

The seventh bit corresponds to bit 6 of Actor_State, which is Defaulted. The system always determines the value of this bit. When bit 6 is set by the system, it indicates that the actor is using defaulted partner information administratively configured for the partner.

The eighth bit corresponds to bit 7 of Actor_State, which is Expired. The system always determines the value of this bit. When bit 7 is set by the system, the actor cannot receive LACPDU frames.

Actor Oper State

The operational state of the port. The Actor Oper State is a string of eight bits that correspond to the operational values of Actor_State, as transmitted by the Actor in Link Aggregation Control Protocol Data Units (LACPDU). The bits are allocated as described for **Actor Admin State** (see above).

Partner Admin State

The administrative state of the partner's port. The Partner Admin State is a string of eight bits that correspond to the administrative value of Actor_State for the protocol Partner.

The first bit corresponds to bit 0 of Actor_State for the Partner, which is Activity. When this bit is set, the dynamic aggregate group is able to exchange LACPDU frames.

The second bit corresponds to bit 1 of Actor_State for the Partner, which is Timeout. When this bit is set, a short timeout is used for LACPDU frames. When this bit is disabled, a long timeout is used for LACPDU frames.

The third bit corresponds to bit 2 of Actor_State for the Partner, which is Aggregation. When this bit is set, the system considers this port to be a potential candidate for aggregation. If this bit is not enabled, the system considers the port to be individual (it can only operate as a single link).

The fourth bit corresponds to bit 3 of Actor_State for the Partner, which is Synchronization. When this bit is set, the port is allocated to the correct dynamic aggregation group. If this bit is not enabled, the port is not allocated to the correct aggregation group.

The fifth bit corresponds to bit 4 of Actor_State for the Partner, which is Collecting. The system always determines the value of this bit. When bit 4 is set by the system, incoming LACPDU frames are collected from the individual ports that make up the dynamic aggregate group.

The sixth bit corresponds to bit 5 of Actor_State for the Partner, which is Distributing. The system always determines the value of this bit. When bit 5 is set by the system, distributing outgoing frames on the port is disabled.

The seventh bit corresponds to bit 6 of Actor_State for the Partner, which is Defaulted. The system always determines the value of this bit. When bit 6 is set by the system, it indicates that the partner is using defaulted actor information administratively configured for the actor.

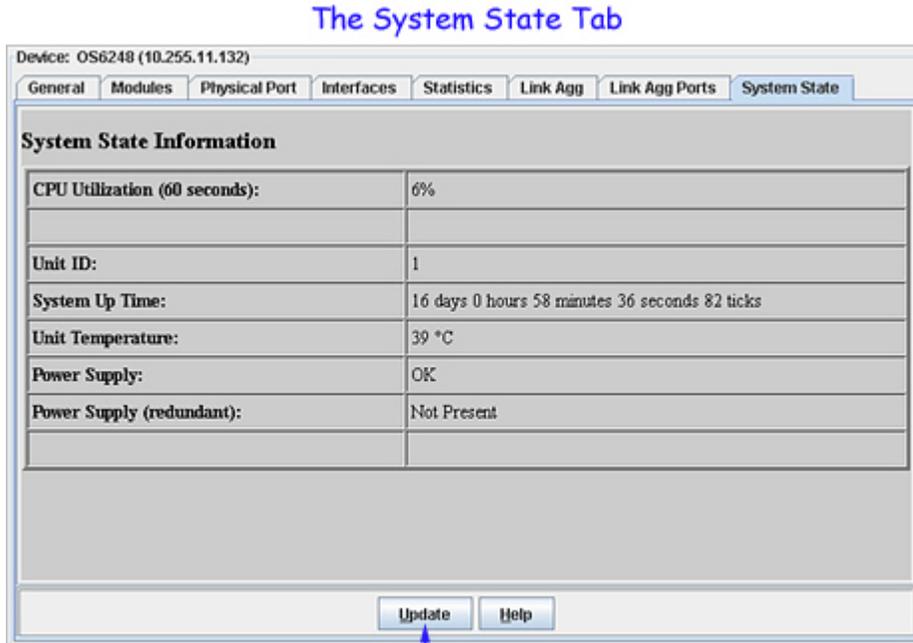
The eighth bit corresponds to bit 7 of Actor_State for the Partner, which is Expired. The system always determines the value of this bit. When bit 7 is set by the system, the partner cannot receive LACPDU frames.

Partner Oper State

The current operational state of the partner's port. The Partner Oper State is a string of eight bits that correspond to the current values of Actor_State in the most recently received Link Aggregation Control Protocol Data Unit (LACPDU) transmitted by the protocol Partner. The bits are allocated as described for **Partner Admin State** (see above).

System State Tab (6200 Devices)

The **System State** tab, displays system state information for each module of the stack.



Click **Update** to poll the switch and refresh the screen with current information.

CPU Utilization (60 seconds)

The average device-level CPU utilization, expressed as a percent, in the primary (active) CMM module over the last 60 seconds.

Unit ID
(TBD)

System Up Time

The time period that has elapsed since the switch was last initialized. (Each tick is .01 second.)

Unit Temperature

This field indicates whether the chassis temperature is within the acceptable temperature range for the switch.

Power Supply

Displays the status of the primary power supply.

Power Supply (redundant)

Displays the status of the redundant power supply.

General Tab (6300-24 Devices)

The General tab for 6300-24 devices provides general system information and general chassis information, as explained below. To change any configurable parameter, edit the field as desired and then click **Apply** to write the change to the switch. All changes take effect immediately.

The General Tab

Device: NMS_125 (10.255.11.125)

General Ports Interfaces

System

Name: NMS_125 Description: OmniStack*24 10/100/1000

Location: NMS LAB Contact:

Up Time: 19 days 4 hours 1 minutes 19 seconds AMAP: Inactive

Chassis Information Table

| Unit Id | Hardware Version | Firmware Version | Loader Version | Boot ROM Version | Op code version | Port Count |
|---------|------------------|------------------|----------------|------------------|-----------------|------------|
| 1 | R01 | | 2.1.0.0 | 2.0.2.5 | 2.2.0.13 | 2 |

Apply Update Help

Click **Update** to poll the switch and refresh the screen with current information.

Click **Apply** to write changes to the switch. All changes take effect immediately.

System Parameters

Name

A user-defined name for this switch.

Description

A factory-defined description of the switch.

Location

A user-defined description of the switch's physical location.

Contact

A user-defined statement identifying the person or organization responsible for the switch.

Up Time

The period of time that has elapsed since the switch was last rebooted.

AMAP

Set this field to **Active** or **Inactive** to enable or disable the AMAP protocol on this switch. By default,

AMAP is enabled. AMAP is a proprietary protocol that learns the connections and links between switches in the list of All Discovered Devices. This information is used to create a graphical display of network links when a network region or subnet is viewed. If you disable AMAP, this switch's connections and links will not be displayed.

Chassis Information Parameters

Unit ID

An ID number that identifies the switch.

Hardware Version

The hardware version of the main board.

Firmware Version

The version of the firmware on the main board.

Loader Version

The version number of the loader code on the main board.

Boot ROM Version

The version number of the Boot ROM and POST (Power On Self Test) code on the main board.

Op Code Version

The version number of the operation (runtime) code on the main board.

Port Count

The total number of ports on the switch, including expansion slots.

Power Status

Indicates whether the switch is using internalPower, redundantPower, or both internalAndRedundantPower.

Expansion Slot 1

The type of module installed in Expansion Slot 1. If no module is installed, **notPresent** displays.

Expansion Slot 2

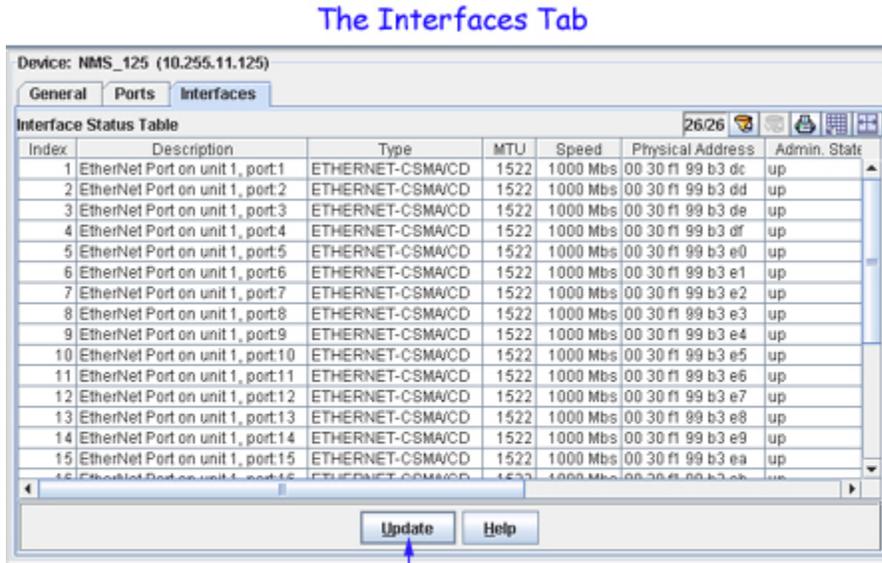
The type of module installed in Expansion Slot 2. If no module is installed, **notPresent** displays.

Role in System

Indicates whether the switch is functioning as the **master**, **backupMaster**, or **slave**.

Interfaces Tab (6300-24 Devices)

The Interfaces tab provides status for all interfaces on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.



Click Update to poll the switch and refresh the screen with current information.

Index

A unique value that identifies this interface.

Description

A textual description of the interface.

Type

The type of the interface, identified according to the physical or link protocol(s) immediately "below" the network layer in the protocol stack.

MTU

The size, in octets, of the largest datagram that can be sent or received on this interface. This is the size of the largest network datagram that can be transmitted on interfaces used for transmitting network datagrams.

Speed

An estimate of the interface's current bandwidth. Speed is displayed in bits-per-second if less than 1,000,000 bits-per-second. Speeds of 1,000,000 bits-per-second or greater are displayed in terms of Mbs (megabits-per-second). If an interface does not vary in bandwidth, or if no accurate estimation can be made, the nominal bandwidth is displayed in this field.

Physical Address

The interface address at the protocol layer (the layer immediately "below" the network layer). This field displays no value for interfaces that do not have such an address (for example, a serial line).

Admin State

The administrative state of the interface: Up indicates the interface is administratively enabled to pass packets; Down indicates the interface is administratively disabled from passing packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Oper Status

The operational state of the interface: Up indicates the interface is able to pass packets; Down indicates the interface is not able to pass packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Last Change

The amount of time since the interface entered its current operational state. This field will display a zero if the current operational state was entered prior to the last re-initialization of the application.

Out Queue

The length of the output packet queue (in packets).

Ports Tab (6300-24 Devices)

The Ports tab provides information on the physical ports on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Ports Tab

| Port Id | Name | Type | Speed Cfg. | Flow Control Cfg. | Auto Negotiation | Speed Status | Flow Co |
|---------|------|---------------|---------------|-------------------|------------------|----------------|---------|
| 1 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 2 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 3 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 4 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 5 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 6 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 7 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 8 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 9 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 10 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 11 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 12 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 13 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 14 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 15 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |
| 16 | | thousandBaseT | fullDuplex100 | disabled | enabled | fullDuplex1000 | none |

Click Update to poll the switch and refresh the screen with current information.

Port ID

An ID number that identifies the port.

Name

The name of the port. This name is the ifAlias in the IF-MIB (RFC2863 or later).

Type

The type of the port.

Speed Cfg.

The speed and duplex mode to which the port is set.

Flow Control Cfg

The flow control mechanism to which the port is set. This field may display:

enabled. Flow control is enabled.

disabled. Flow control is disabled.

backPressure. Flow control mechanism is backPressure when the port is in fullDuplex mode. This flow control mechanism will not function.

dot3xFlowControl. Flow control mechanism is IEEE 802.3x flow control when the port is in halfDuplex mode. This flow control mechanism will not function.

Auto Negotiation

The status of auto negotiation: **enabled** or **disabled**.

Speed Status

The speed and duplex mode at which the port is operating. If this port is operating as a trunk, the speed is the speed of its individual members. If this port is operating as a trunk and the result is inconsistent among its member ports, this field will display **error**.

Flow Control Status

The flow control mechanism that the port is actually using. This field may display:

error. This is a trunk and the result is inconsistent among its member ports.

backPressure. The BackPressure flow control mechanism is being used.

dot3xFlowControl. The IEEE 802.3 flow control mechanism is being used.

none. Flow control is disabled.

Forced Mode

The forced mode of a combination port (ports 21 - 24). If this port is not a combination port, this field displays **none**. If the port is a combination port, this field may display:

copperForced. Always uses the built-in RJ-45 port.

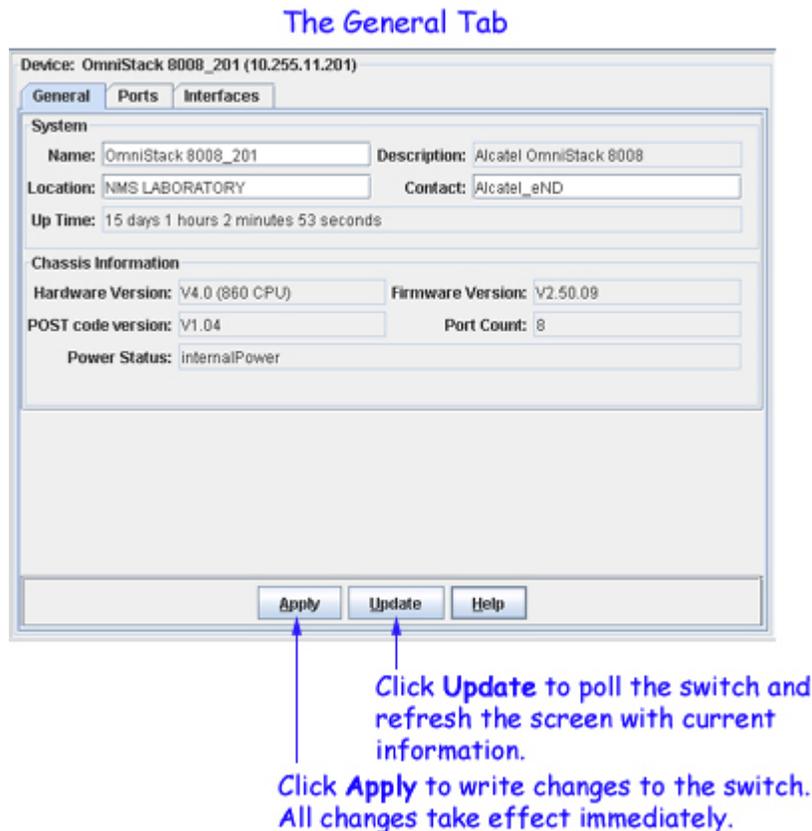
copperPreferredAuto. Uses the built-in RJ-45 port if both combination types are functioning and if the RJ-45 port has a valid link

sfpForced. Always uses the SFP port (even if the module is not installed)

sfpPreferredAuto. Uses the SFP port if both combination types are functioning and if the SFP port has a valid link

General Tab (8008 Devices)

The General tab for 8008 devices provides general system information and general chassis information, as explained below. To change any configurable parameter, edit the field as desired and then click **Apply** to write the change to the switch. All changes take effect immediately.



System Parameters

Name

A user-defined name for this switch.

Description

A description of the switch as defined by the manufacturer.

Location

A user-defined description of the switch's physical location.

Contact

A user-defined parameter stating who is responsible for this switch.

Up Time

The period of time that has elapsed since the switch was last rebooted.

Chassis Information Parameters

Hardware Version

The version number of the main hardware board.

Firmware Version

The version number of the system firmware in flash ROM.

POST Code Version

The version number of the POST (Power-On Self-Test) code in ROM.

Port Count

The total number of ports on the switch.

Power Status

Displays the type of power the switch is using: `internalPower`, `redundantPower`, or `internalAndRedundantPower`.

Interfaces Tab (8008 Devices)

The Interfaces tab provides status for all interfaces on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Interfaces Tab

Device: OmniStack 8008_201 (10.255.11.201)

General Ports Interfaces

Interface Status Table

| Index | Description | Type | MTU | Speed | Physical Address | Admin |
|-------|-----------------------|------------------|------|----------|-------------------|-------|
| 1 | RMON Port 1 on Unit 1 | ETHERNET-CSMA/CD | 1500 | 1000 Mbs | 00 d0 95 4b 07 21 | up |
| 2 | RMON Port 2 on Unit 1 | ETHERNET-CSMA/CD | 1500 | 1000 Mbs | 00 d0 95 4b 07 22 | up |
| 3 | RMON Port 3 on Unit 1 | ETHERNET-CSMA/CD | 1500 | 1000 Mbs | 00 d0 95 4b 07 23 | up |
| 4 | RMON Port 4 on Unit 1 | ETHERNET-CSMA/CD | 1500 | 1000 Mbs | 00 d0 95 4b 07 24 | up |
| 5 | RMON Port 5 on Unit 1 | ETHERNET-CSMA/CD | 1500 | 1000 Mbs | 00 d0 95 4b 07 25 | up |
| 6 | RMON Port 6 on Unit 1 | ETHERNET-CSMA/CD | 1500 | 1000 Mbs | 00 d0 95 4b 07 26 | up |
| 7 | RMON Port 7 on Unit 1 | ETHERNET-CSMA/CD | 1500 | 1000 Mbs | 00 d0 95 4b 07 27 | up |
| 8 | RMON Port 8 on Unit 1 | ETHERNET-CSMA/CD | 1500 | 1000 Mbs | 00 d0 95 4b 07 28 | up |
| 1001 | Console port | 33 | 1500 | 0 | | up |
| 1101 | Management Port | OTHER | 1500 | 10 Mbs | 00 d0 95 4b 07 20 | up |

Update Help

Click Update to poll the switch and refresh the screen with current information.

Index

A unique value that identifies this interface.

Description

A textual description of the interface.

Type

The type of the interface, identified according to the physical or link protocol(s) immediately "below" the network layer in the protocol stack.

MTU

The size, in octets, of the largest datagram that can be sent or received on this interface. This is the size of the largest network datagram that can be transmitted on interfaces used for transmitting network datagrams.

Speed

An estimate of the interface's current bandwidth. Speed is displayed in bits-per-second if less than 1,000,000 bits-per-second. Speeds of 1,000,000 bits-per-second or greater are displayed in terms of Mbs (megabits-per-second). If an interface does not vary in bandwidth, or if no accurate estimation can be made, the nominal bandwidth is displayed in this field.

Physical Address

The interface address at the protocol layer (the layer immediately "below" the network layer). This field displays no value for interfaces that do not have such an address (for example, a serial line).

Admin State

The administrative state of the interface: Up indicates the interface is administratively enabled to pass packets; Down indicates the interface is administratively disabled from passing packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Oper Status

The operational state of the interface: Up indicates the interface is able to pass packets; Down indicates the interface is not able to pass packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Last Change

The amount of time since the interface entered its current operational state. This field will display a zero if the current operational state was entered prior to the last re-initialization of the application.

Out Queue

The length of the output packet queue (in packets).

Ports Tab (8008 Devices)

The Ports tab provides information on the physical ports on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Ports Tab

Click **Update** to poll the switch and refresh the screen with current information.

| Port Id | Port Type | Admin Speed and Mode | Oper Speed and Mode | Admin Flow Control |
|---------|----------------|----------------------|---------------------|--------------------|
| 1 | thousandBaseSX | autoNegotiation | fullDuplex1000 | disabled |
| 2 | thousandBaseSX | autoNegotiation | fullDuplex1000 | disabled |
| 3 | thousandBaseSX | autoNegotiation | fullDuplex1000 | disabled |
| 4 | thousandBaseSX | autoNegotiation | fullDuplex1000 | disabled |
| 5 | thousandBaseSX | autoNegotiation | fullDuplex1000 | disabled |
| 6 | thousandBaseSX | autoNegotiation | fullDuplex1000 | disabled |
| 7 | thousandBaseSX | autoNegotiation | fullDuplex1000 | disabled |
| 8 | thousandBaseSX | autoNegotiation | fullDuplex1000 | disabled |

Port ID

An ID number that identifies the port within this switch.

Port Type

The type of the port.

Admin Speed and Mode

The speed and duplex mode to which the port is set administratively. The value in this field may be **halfDuplex1000** (1000 Mbps and half duplex mode), **fullDuplex1000** (1000 Mbps and full duplex mode), or **autoNegotiation** (allow the switch to negotiate duplex mode and speed with the other end of connection).

Oper Speed and Mode

The speed and duplex mode at which the port is actually operating. The value in this field may be **halfDuplex1000** or **fullDuplex1000**.

Admin Flow Control

The administrative state of flow control for the port: either **enabled** or **disabled**. When flow control is enabled, and the port is operating in halfDuplex mode, the backPressure flow control mechanism is used. When flow control is enabled, and the port is operating in fullDuplex mode, the IEEE 802.3x flow control mechanism is used. Flow control can eliminate frame loss by “blocking” traffic from end stations or segments connected directly to the switch when switch buffers fill.

Oper Flow Control

The type of flow control the port is actually using during operation. This field may display the following values:

backPressure. The backPressure flow control mechanism is in use. The backPressure flow control mechanism is used when flow control is administratively enabled and the port is operating in halfDuplex mode at 1000 Mbps.

dot3xFlowControl. The IEEE 802.3x flow control mechanism is in use. The IEEE 802.3x flow control mechanism is used when flow control is administratively enabled and the port is operating in fullDuplex mode at 1000 Mbps.

none. Flow control is disabled.

General Tab (Third-Party Devices)

The General tab for third-party devices enables you to specify the device name and location of the third-party device. It also displays the system up time (the period of time that has elapsed since the switch was last rebooted). To change the device name or location, edit the respective fields as desired and then click **Apply** to write the change to the switch. All changes take effect immediately.

The General Tab

The screenshot shows a web interface for configuring a third-party device. At the top, it says "Device: OAW-4308 (10.255.11.151)". Below this are two tabs: "General" (selected) and "Interfaces". The "System" section contains several fields: "Name" (OAW-4308), "Description" (AOS-W (MODEL: OAW-4308), Version 2.4.2.0), "Location" (public-public md5+des user w 'public' trap use), "Contact" (configure ov with 'public' trap user for v2 traps), and "Up Time" (4 days 7 hours 5 minutes 21 seconds). At the bottom of the form are three buttons: "Apply", "Update", and "Help".

Click **Update** to poll the switch and refresh the screen with current information.

Click **Apply** to write changes to the switch. All changes take effect immediately.

Interfaces Tab (Third-Party Devices)

The Interfaces tab provides status for all interfaces on the switch. Click once in any column header to display the Down Arrow and sort table information in ascending order. Click a second time to display the Up Arrow and sort in descending order. Each field is described below.

The Interfaces Tab

Device: OAW-4308 (10.255.11.151)

General Interfaces

Interface Status Table 11/11

| Index | Description | Type | MTU | Speed | Physical Address | Admin. State |
|-----------|---------------------|--------------------|------|-----------|-------------------|--------------|
| 4097 | fe1/0 | ETHERNET-CRMA/CD | 1500 | 100 Mbps | 00 0b 86 50 6f e6 | up |
| 4098 | fe1/1 | ETHERNET-CRMA/CD | 1500 | 100 Mbps | 00 0b 86 50 6f e7 | up |
| 4099 | fe1/2 | ETHERNET-CRMA/CD | 1500 | 100 Mbps | 00 0b 86 50 6f e8 | up |
| 4100 | fe1/3 | ETHERNET-CRMA/CD | 1500 | 100 Mbps | 00 0b 86 50 6f e9 | up |
| 4101 | fe1/4 | ETHERNET-CRMA/CD | 1500 | 100 Mbps | 00 0b 86 50 6f ea | up |
| 4102 | fe1/5 | ETHERNET-CRMA/CD | 1500 | 100 Mbps | 00 0b 86 50 6f eb | up |
| 4103 | fe1/6 | ETHERNET-CRMA/CD | 1500 | 100 Mbps | 00 0b 86 50 6f ec | up |
| 4104 | fe1/7 | ETHERNET-CRMA/CD | 1500 | 100 Mbps | 00 0b 86 50 6f ed | up |
| 4105 | gig1/8 | ETHERNET-CRMA/CD | 1500 | 1000 Mbps | 00 0b 86 50 6f ee | up |
| 16395 | 802.1Q VLAN | L3 VLAN (IP) | 1500 | | 00 0b 86 50 6f e5 | up |
| 134217728 | SWITCH IP INTERFACE | SOFTWARE-LOOP-BACK | 1500 | | 00 00 00 00 00 00 | up |

Update Help

Click Update to poll the switch and refresh the screen with current information.

Index

A unique value that identifies this interface.

Description

A textual description of the interface.

Type

The type of the interface, identified according to the physical or link protocol(s) immediately "below" the network layer in the protocol stack.

MTU

The size, in octets, of the largest datagram that can be sent or received on this interface. This is the size of the largest network datagram that can be transmitted on interfaces used for transmitting network datagrams.

Speed

An estimate of the interface's current bandwidth. Speed is displayed in bits-per-second if less than 1,000,000 bits-per-second. Speeds of 1,000,000 bits-per-second or greater are displayed in terms of Mbs (megabits-per-second). If an interface does not vary in bandwidth, or if no accurate estimation can be made, the nominal bandwidth is displayed in this field.

Physical Address

The interface address at the protocol layer (the layer immediately "below" the network layer). This field displays no value for interfaces that do not have such an address (for example, a serial line).

Admin. State

The administrative state of the interface: Up indicates the interface is administratively enabled to pass packets; Down indicates the interface is administratively disabled from passing packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

Oper. Status

The operational state of the interface: Up indicates the interface is able to pass packets; Down indicates the interface is not able to pass packets; Testing indicates the interface is in a test mode and no operational packets can be passed.

LastChange

The amount of time since the interface entered its current operational state. This field will display a zero if the current operational state was entered prior to the last reinitialization of the application.

OutQueue

The length of the output packet queue (in packets).

Importing MIBs

The **Import MIBs** menu item on the File menu, shown below, enables you to import new or updated MIB files into OmniVista. All MIB files are imported to the OmniVista server.

Before You Begin

Before you import MIBs, it is important to understand that the the purpose of this function is to import MIB files that reside somewhere on your local file system into OmniVista. The end result of this operation is that the imported MIBs will reside in the *installationroot/data/mibs* directory on the server. A *mibs.txt* ASCII file lists the order in which the MIBs will be compiled. It is NOT recommended that you manually copy MIB files that you want to import into the *installationroot/data/mibs* directory.

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 also.

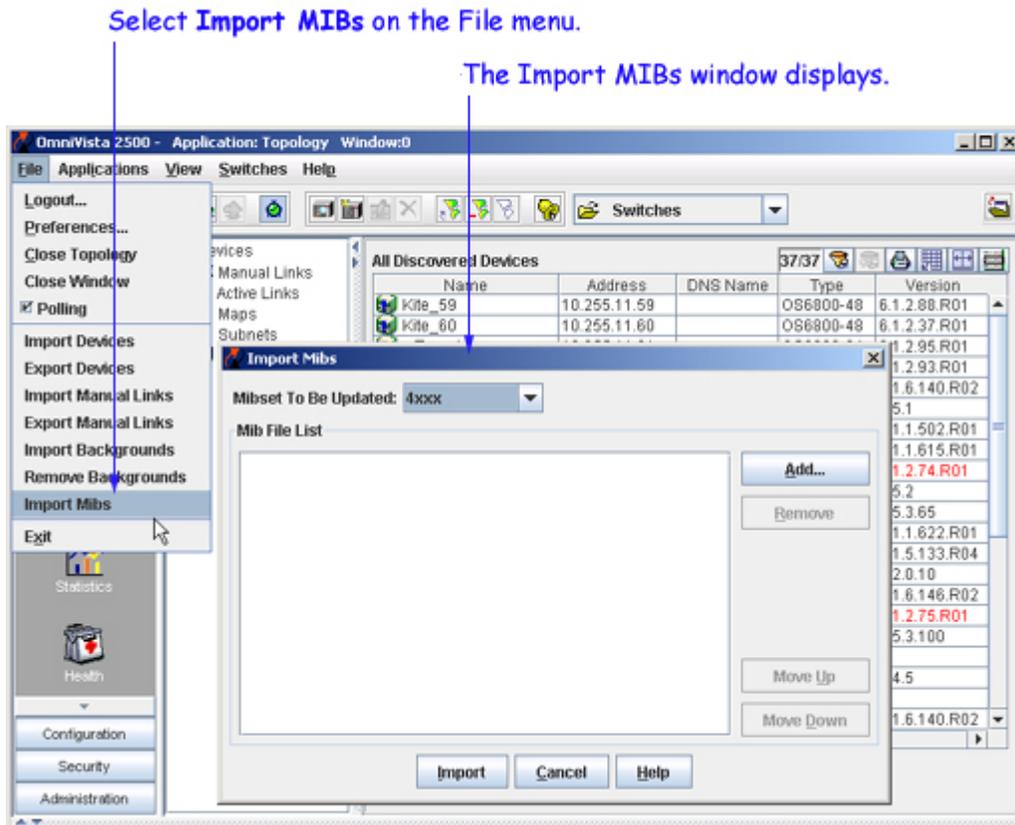
In order for the MIBs to compile correctly, you are strongly advised to order them so that all the referenced MIB files are compiled before the files that reference them. MIB compilers follow import references from one MIB to another on the fly, and do not strictly require that the MIBs be compiled in 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 using the **Move Up** and **Move Down** buttons in the Import MIBs window, as shown and described in the procedure below. MIB files will be compiled in the order that the files are listed in the Import MIBs window.

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.

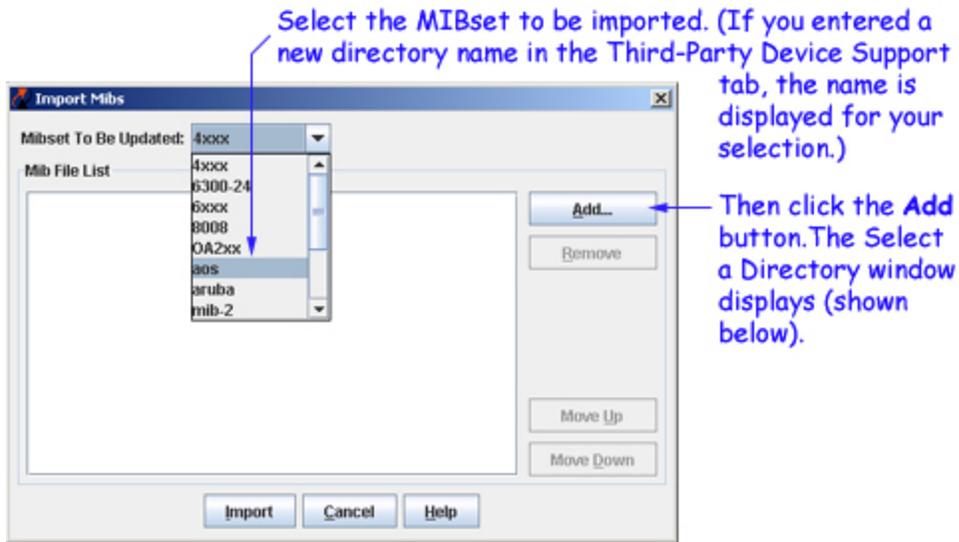
Once you have completed the MIB importation process, OmniVista does not immediately parse the MIBs. When you discover a device with an OID that is specified for the MIB directory into which you imported the new MIBs, OmniVista will poll the device for standard MIB-II objects. If the standard MIB-II MIBs 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.

How to Import MIBs

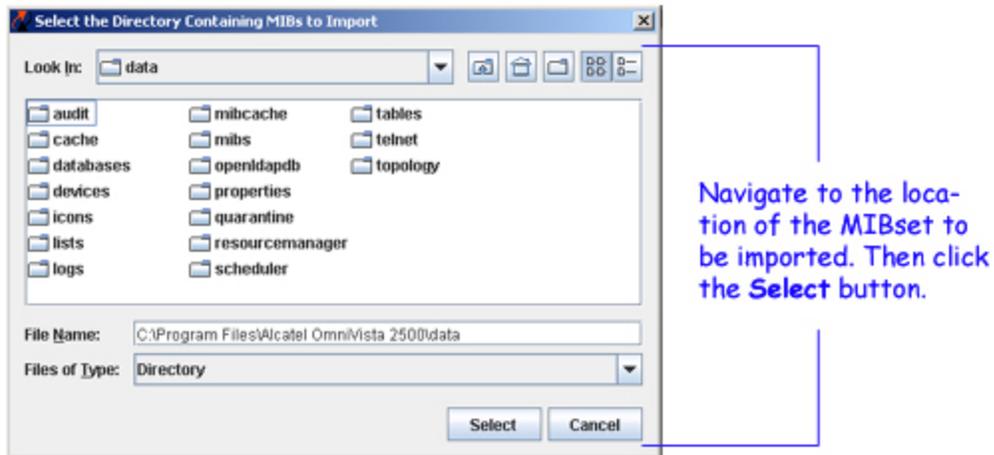
1. Execute the Topology application and select **Import MIBs** on the File menu. The Import MIBs window displays, shown below.



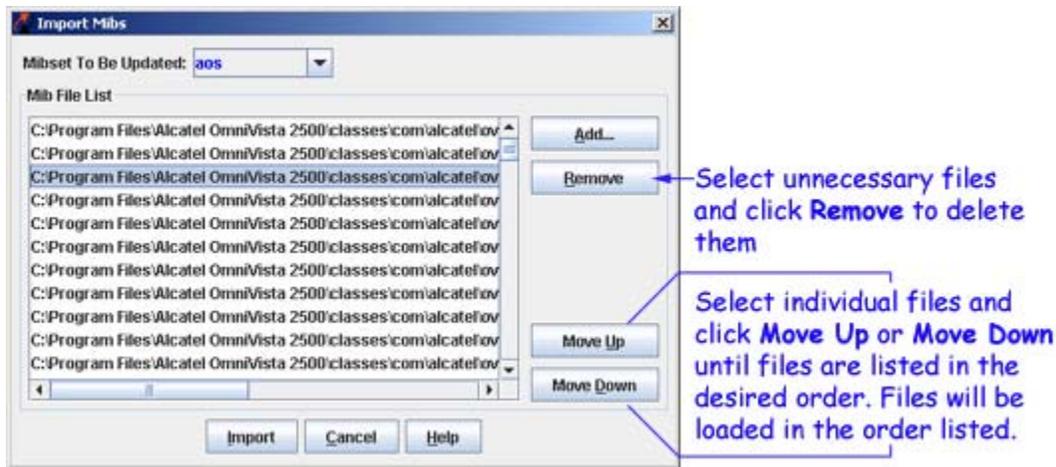
2. Set the MIBset to be Updated combo box (shown below) to the MIBset to be imported. (If you entered a new directory name in the Third-Party Device Support tab, the name is displayed for your selection.) Then click the **Add** button. The Select a Directory window displays.



The Select a Directory Window



3. Within the Select a Directory window, navigate to the location where the MIBset resides. When the correct MIB directory is displayed in the window, click the **Select** button. The Select a Directory window closes and the MIB files are listed in the Import MIBs window, as shown below.



4. If any files listed in the Import MIBs window are unnecessary, select them and click the **Remove** button. Files that you remove will not be imported.

5. The MIB files will be loaded into OmniVista in the order the files are listed in the the Import MIBs window. You can adjust this order by selecting individual files and clicking the **Move Up** and **Move Down** buttons until files are listed in the correct order.

6. Click the **Import** button. The MIB files are imported to the OmniVista server. A message displays in the Status Panel when the import operation is complete.