Configuration — VLANs, Spanning Tree, and Multi-Link Trunking
Avaya Ethernet Routing Switch 2500 Series

4.4
NN47215-501, 05.03
August 2011

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

“Documentation” means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya generally makes available to users of its products. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked Web sites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on its Hardware and Software (“Product(s)”). Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya’s standard warranty language, as well as information regarding support for this Product while under warranty is available to Avaya customers and other parties through the Avaya Support Web site: http://support.avaya.com. Please note that if you acquired the Product(s) from an authorized Avaya reseller outside of the United States and Canada, the warranty is provided to you by said Avaya reseller and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO/ ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA AUTHORIZED RESELLER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE (“AVAYA”).

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, or Hardware provided by Avaya. All content on this site, the documentation and the Product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third-party components

Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements (“Third Party Components”), which may contain terms that expand or limit rights to use certain portions of the Product (“Third Party Terms”). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: http://support.avaya.com/Copyright.

Trademarks

The trademarks, logos and service marks (“Marks”) displayed in this site, the Documentation and Product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation and Product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners, and “Linux” is a registered trademark of Linus Torvalds.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support Web site: http://support.avaya.com.

Contact Avaya Support

Avaya provides a telephone number for you to use to report problems or to ask questions about your Product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: http://support.avaya.com.

Contents

Chapter 1: New in this release

The following section detail the new features in Configuration - VLANs, Spanning Tree, and Multi-Link Trunking (NN47215-501) for Release 4.4.

ADAC enhancements

Auto-Detect Auto-Configuration (ADAC) enhancements provide increased flexibility in deployments that use ADAC as follows:

- Expanded support for up to eight ADAC uplinks and eight Call Server links (individual ports or any combination of MLT, DMLT, or LAG) for each switch or stack.
- Ability to change the non-ADAC VLANs on a port without disabling ADAC.

IPv6 protocol-based VLANs

IPv6 recognition through the configuration of protocol-based VLANs for segmenting IPv6 traffic is supported.

MLT enable or disable whole trunk

Use the Multi-Link Trunk (MLT) enable or disable whole trunk feature to enable or disable trunk loop prevention for MLT or Distributed MLT (DMLT). The feature is disabled by default. If you enable the feature, the state of the port changes to reflect the state of the MLT or DMLT bundle irrespective of the previous status. If you disable the MLT or DMLT then all links that are part of the MLT group are disabled, with the exception of the Destination Lookup Failure (DLF) link. For network configuration, Avaya recommends you to enable the MLT whole trunk feature.

Chapter 2: Introduction

This document provides information you need to configure VLANs, Spanning Tree and Multi-Link Trunking for the Ethernet Routing Switch 2500 Series.

ACLI command modes

ACLI provides the following command modes:

- User EXEC
- Privileged EXEC
- Global Configuration
- Interface Configuration

Mode access is determined by access permission levels and password protection.

If no password is set, you can enter ACLI in User EXEC mode and use the **`enable`** command to move to the next level (Privileged EXEC mode). However, if you have read-only access, you cannot progress beyond User EXEC mode, the default mode. If you have read-write access you can progress from the default mode through all of the available modes.

With sufficient permission, you can use the rules in the following table to move between the command modes.

Command mode and sample prompt	Entrance commands	Exit commands

User EXEC
`2526T>`
	No entrance command, default mode	`exit`
		or
		`logout`

Privileged EXEC
`2526T#`
	`enable`	`exit`
		or
		`logout`

Global Configuration
`2526T(config)#`
	`configure`	To return to Privileged EXEC mode, enter:
		`end`
		or
		`exit`
		To exit ACLI completely, enter:

Command mode and sample prompt	Entrance commands	Exit commands
		`logout`

Interface Configuration
`2526T(config-if)#`
	From Global Configuration mode: To configure a port, enter: interface fastethernet <port number> To configure a VLAN, enter: interface vlan <vlan number>	To return to Global Configuration mode, enter:
		`exit`
		To return to Privileged EXEC mode, enter:
		`end`
		To exit ACLI completely, enter:
		`logout`

See *Avaya Ethernet Routing Switch 2500 Series Fundamentals* (NN47215-102).

Chapter 3: VLAN Fundamentals

Virtual local area networks

In a traditional shared-media network, traffic that a station generates is transmitted to all other stations on the local segment. Therefore, for any given station on the shared Ethernet, the local segment is the collision domain because traffic on the segment has the potential to cause an Ethernet collision. The local segment is also the *broadcast domain* because any broadcast is sent to all stations on the local segment. Although Ethernet Routing Switches and bridges divide a network into smaller collision domains, they do not affect the broadcast domain. In simple terms, a virtual local area network (VLAN) provides a mechanism to fine-tune broadcast domains.

With the Ethernet Routing Switch 2500 Series , you can to create port-based and IPv6 protocol-based virtual local area networks (VLANs):

- IEEE 802.1Q port-based VLANs

  A port-based VLAN is a VLAN in which the switch ports are explicitly configured to be in the VLAN. When you create a port-based VLAN, you assign a Port VLAN Identifier (PVID) and specify which ports belong to the VLAN. The PVID is used to coordinate VLANs across multiple switches.

- IPv6 protocol-based VLANs

  A protocol-based VLAN is a VLAN in which the switch examines the protocol in use on the port. When you create a protocol-based VLAN, you assign a protocol ID for the VLAN. IPv6 recognition for segmenting IPv6 traffic is supported.

- VLAN Configuration Control

  VLAN Configuration Control (VCC) to modify VLANs. VLAN Configuration Control is a superset of the existing AutoPVID functionality and incorporates this functionality for backward compatibility. VLAN Configuration Control is globally applied to all VLANs on the switch.

  For more information, see VLAN Configuration Control on page 18.

VLAN support

The Ethernet Routing Switch 2500 Series supports 256 VLANs, either by-port, under the 802.1d bridging model, or IPv6 protocol-based VLANs.

PVIDs are by port assignment. The AutoPVID option automatically assigns a PVID to all the ports. These ports are the members of the VLAN that are created.

When the Ethernet Routing Switch 2500 Series is installed for the first time, all ports are assigned to the default VLAN (PVID = 1). The default management VLAN is VLAN 1.

You can configure VLANs through the ACLI or EDM interfaces. The Ethernet Routing Switch 2500 Series supports binary and ASCII configuration files. You can also configure VLANs using both SNMP and ASCII scripts.

IEEE 802.1Q tagging

The Ethernet Routing Switch 2500 Series allows tagging by port on all ports. Tagging status applies on all ports of a Multi-Link trunk (a port member in a Multi-Link trunk cannot be configured independently of the other members in the same Multi-Link trunk). You can configure untagged frame dropping by port.

Ethernet Routing Switch 2500 Series supports the Independent VLAN Learning (IVL) model. IVL allows duplicate MAC address to be present in different sets, but not in the same set or VLAN.

IEEE 802.1Q VLAN workgroups

The Ethernet Routing Switch 2500 Series supports up to 256 VLANs and the Ethernet Routing Switch 2500 Series supports IEEE 802.1Q tagging available for each per port. Ports are grouped into broadcast domains by assigning them to the same VLAN. Frames received in one VLAN can only be forwarded within that VLAN, and multicast frames and unknown unicast frames are flooded only to ports in the same VLAN.

When you set up VLANs, you segment networks to increase network capacity and performance without changing the physical network topology (Figure 1: Port-based VLAN example on page 15). With network segmentation, each switch port connects to a segment that is a single broadcast domain. When a switch port is configured to be a member of a VLAN, it is added to a group of ports (workgroup) that belong to one broadcast domain.

You can use the Ethernet Routing Switch 2500 Series to assign ports to VLANs using the console, Telnet or an appropriate SNMP-based application. You can assign different ports (and therefore the devices attached to these ports) to different broadcast domains. This feature allows network flexibility because you can reassign VLANs to accommodate network moves, additions, and changes, eliminating the need to change physical cabling.

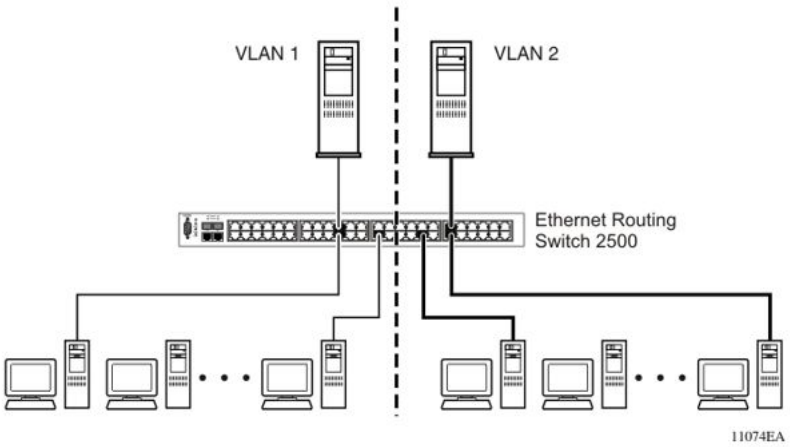

Figure 1: Port-based VLAN example

IEEE 802.1Q tagging

The Ethernet Routing Switch 2500 Series operates in accordance with the IEEE 802.1Q tagging rules. Important terms used with the 802.1Q tagging feature are:

- VLAN identifier (VID)—the 12-bit portion of the VLAN tag in the frame header that identifies an explicit VLAN.
- Port VLAN identifier (PVID)—a classification mechanism that associates a port with a specific VLAN. For example, a port with a PVID of 3 (PVID =3) assigns all untagged frames received on this port to VLAN 3.
- Tagged frame—the 32-bit field (VLAN tag) in the frame header that identifies the frame as belonging to a specific VLAN. Untagged frames are marked (tagged) with this classification as they leave the switch through a port that is configured as a tagged port.
- Untagged frame—a frame that does not carry any VLAN tagging information in the frame header.
- VLAN port members—a set of ports that form a broadcast domain for a specific VLAN. A port can be a member of one or more VLANs.
- Untagged member—a port that is configured as an untagged member of a specific VLAN. When an untagged frame exits the switch through an untagged member port, the frame header remains unchanged. When a tagged frame exits the switch through an untagged member port, the tag is stripped and the tagged frame is changed to an untagged frame.
- Tagged member—a port that is configured as a member of a specific VLAN. When an untagged frame exits the switch through a tagged member port, the frame header is modified to include the 32-bit tag associated with the PVID. When a tagged frame exits the switch through a tagged member port, the frame header remains unchanged (original VID remains).

- User priority—a three-bit field in the header of a tagged frame. The field is interpreted as a binary number, and therefore has a value of 0 to 7. This field allows the tagged frame to carry the user priority across bridged LANs in which the individual LAN segments are sometimes unable to signal priority information.
- Port priority—the priority level assigned to *untagged* frames received on a port. This value becomes the user priority for the frame. *Tagged* packets get their user priority from the value contained in the 802.1Q frame header.
- Unregistered packet—a tagged frame that contains a VID where the receiving port is not a member of that VLAN.

By default, all Ethernet Routing Switch 2500 Series ports are set as untagged members of VLAN 1 with all ports configured as PVID = 1. Every VLAN is assigned a unique VID that distinguishes it from all other VLANs. In the default configuration example shown in Figure 2: Default VLAN settings on page 16, all incoming packets are assigned to VLAN 1 by the default port VLAN identifier (PVID =1). Untagged packets enter and leave the switch unchanged.

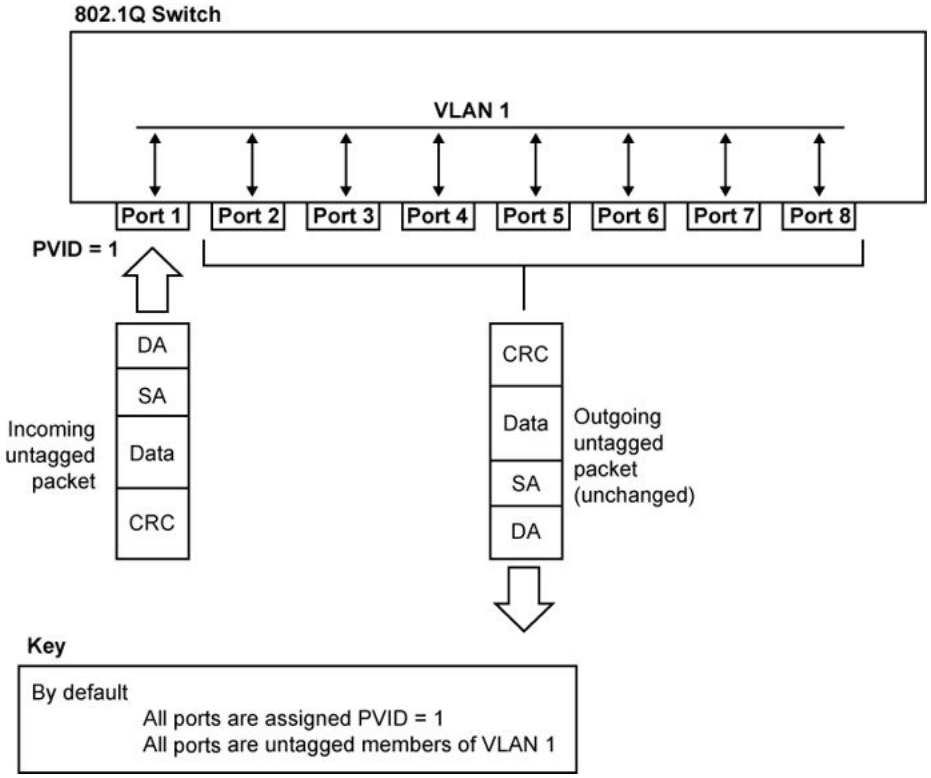

Figure 2: Default VLAN settings

When you configure VLANs, you configure the switch ports as *tagged* or *untagged* members of specific VLANs (see Figure 3: Port-based VLAN assignment on page 17 through Figure 11: VLAN broadcast domains within the switch on page 23).

In Figure 3: Port-based VLAN assignment on page 17, untagged incoming packets are assigned directly to VLAN 2 (PVID = 2). Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.

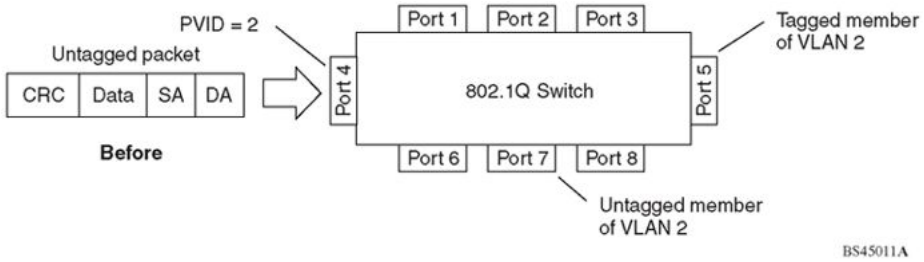

Figure 3: Port-based VLAN assignment

As shown in Figure 4: 802.1Q tagging (after port-based VLAN assignment) on page 17, the untagged packet is marked (tagged) as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. The untagged packet remains unchanged as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

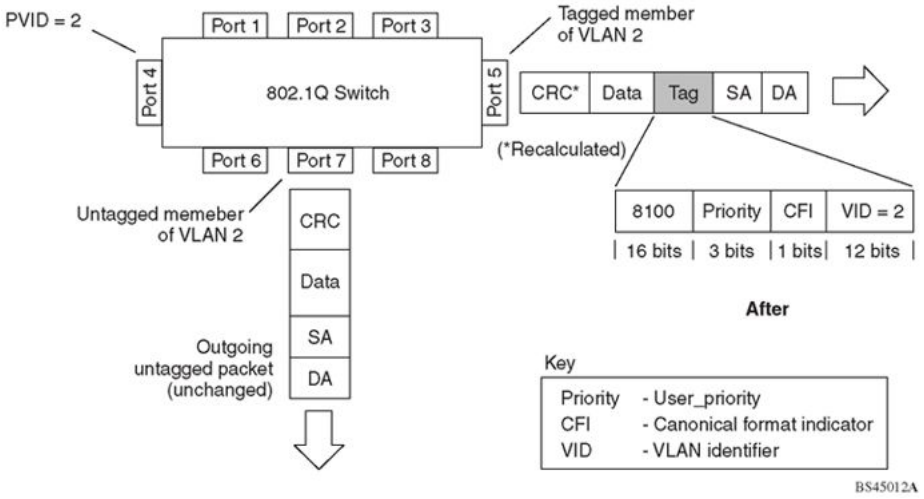

Figure 4: 802.1Q tagging (after port-based VLAN assignment)

In Figure 5: 802.1Q tag assignment on page 17, tagged incoming packets are assigned directly to VLAN 2 because of the tag assignment in the packet. Port 5 is configured as a tagged member of VLAN 2, and port 7 is configured as an untagged member of VLAN 2.

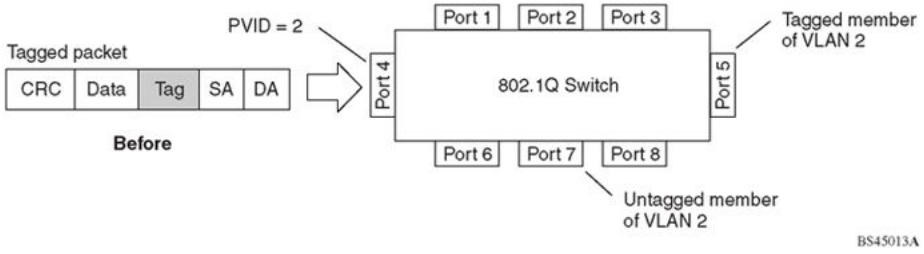

Figure 5: 802.1Q tag assignment

As shown in Figure 6: 802.1Q tagging (after 802.1Q tag assignment) on page 18, the tagged packet remains unchanged as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. However, the tagged packet is stripped (untagged) as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

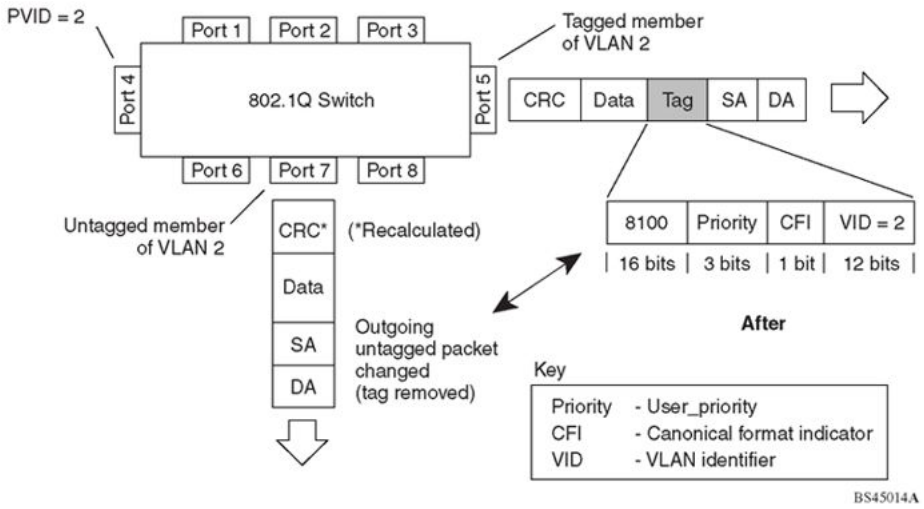

Figure 6: 802.1Q tagging (after 802.1Q tag assignment)

VLAN Tagging Enhancement

Release 4.2 or later provides additional options for VLAN port tagging. Rather than setting a port to untagged or tagged mode, you can also choose to enable or disable PVID tagging.

Table 1: VLAN Tagging mode definitions on page 18 summarizes the new tagging options.

Table 1: VLAN Tagging mode definitions

Tagging mode	Definition	
	PVID Tagging	Non-PVID Tagging
Untag All (Untagged Access)	Disabled	Disabled
Tag All (Tagged Trunk)	Enabled	Enabled
Tag PVID Only	Enabled	Disabled
Untag PVID Only	Disabled	Enabled

VLAN Configuration Control

Switch administrators use VLAN Configuration Control (VCC) to control how VLANs are modified. VLAN Configuration Control is a superset of the existing AutoPVID functionality and incorporates this functionality for backward compatibility. VLAN Configuration Control is globally applied to all VLANs on the switch.

VLAN Configuration Control offers four options for controlling VLAN modification:

1. **Strict**—This option restricts the addition of an untagged port to a VLAN if the port is already a member of another VLAN. To add an untagged port to a new VLAN,

the switch administrator must remove the port from all other VLANs of which it is a member of before adding it to the new VLAN. The PVID of the port will be changed to the new VID to which it was added.

Important:

Strict is the factory default setting.

2. **Automatic**—This option automatically adds an untagged port to a new VLAN and automatically removes it from any previous VLAN membership. The PVID of the port is automatically changed to the VID of the VLAN it joins. Because the port is first added to the new VLAN and then removed from any previous membership, the Spanning Tree Group participation of the port will not be disabled as long as the VLANs involved are in the same Spanning Tree Group.
3. **AutoPVID**—This option functions in the same manner as previous AutoPVID functionality. When an untagged port is added to a new VLAN, the port is added to the new VLAN and the PVID assigned to the new VID without removing it from any previous VLAN memberships. When using this option, an untagged port has membership in multiple VLANs.
4. **Flexible**—This option functions in a similar manner to disabling AutoPVID functionality. When this option is used, there are no restrictions on the number of VLANs to which an untagged port can belong. Any new additions of an untagged port to a new VLAN does not change the PVID of that port.

VLAN Configuration Control is only applied to ports with the tagging modes of Untag All and Tag PVID Only. VLAN Configuration Control does not control ports with the tagging modes of Tag All and Untag PVID Only. Ports with the tagging modes of Tag All and Untag PVID Only can belong to multiple VLANs regardless of VLAN Configuration Control settings and their PVID must be manually changed.

VLANs spanning multiple switches

You can use VLANs to segment a network within a switch. When you connect multiple switches, it is possible to connect users of one VLAN with users of the same VLAN in another switch. However, the configuration guidelines depend on whether both switches support 802.1Q tagging.

With 802.1Q tagging enabled on a port for a VLAN, all frames leaving the port for that VLAN are marked as belonging to that specific VLAN. You can assign specific switch ports as members of one or more VLANs that span multiple switches, without interfering with the Spanning Tree Protocol.

VLANs spanning multiple 802.1Q tagged switches

The following figure shows VLANs spanning two Ethernet Routing Switch 2500 Series devices. The 802.1Q tagging is enabled on S1, port 2 and on S2, port 1 for VLAN 1 and VLAN 2. Both ports are tagged members of VLAN 1 and VLAN 2.

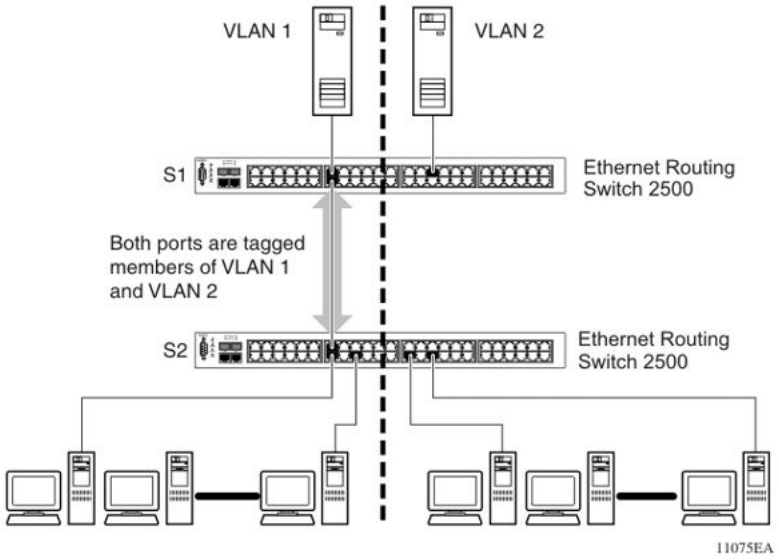

Figure 7: VLANs spanning multiple 802.1Q tagged switches

Because there is only one link between the two switches, the Spanning Tree Protocol (STP) treats this configuration as any other switch-to-switch connection. For this configuration to work properly, both switches must support the 802.1Q tagging protocol.

VLANs spanning multiple untagged switches

Figure 8: VLANs spanning multiple untagged switches on page 21 shows VLANs spanning multiple untagged switches. In this configuration, S2 does not support 802.1Q tagging and you must use a single switch port on each switch for each VLAN.

For this configuration to work properly, you must set Spanning Tree participation to Disabled (the STP is not supported across multiple LANs).

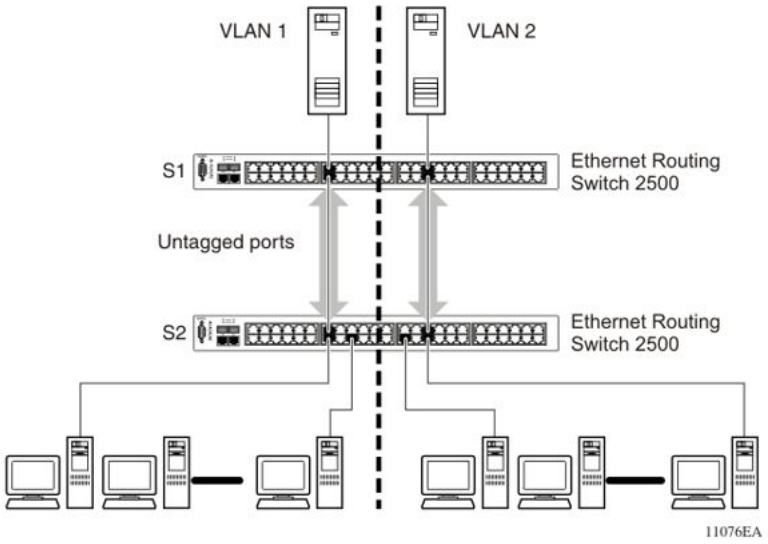

Figure 8: VLANs spanning multiple untagged switches

When the STP is enabled on these switches, only one link between each pair of switches forwards traffic. Because each port belongs to only one VLAN at a time, connectivity on the other VLAN is lost. Exercise care when configuring the switches to ensure that the VLAN configuration does not conflict with spanning tree configuration.

To connect multiple VLANs across switches with redundant links, you must disable the STP on all participating switch ports. Figure 9: Possible problems with VLANs and Spanning Tree Protocol on page 21 shows possible consequences of enabling the STP when using VLANs between untagged (non-802.1Q tagged) switches.

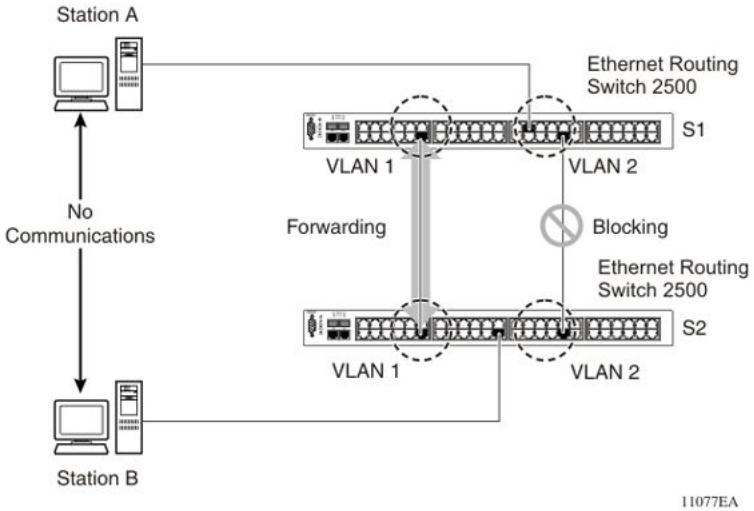

Figure 9: Possible problems with VLANs and Spanning Tree Protocol

As shown in Figure 9: Possible problems with VLANs and Spanning Tree Protocol on page 21, with STP enabled, only one connection between S1 and S2 is forwarding at any time.

Communications failure occurs between VLAN 2 of S1 and VLAN 2 of S2, blocking communications between Stations A and B.

The STP selects the link connecting VLAN 1 on S1 and S2 as the forwarding link based on port speed, duplex-mode, and port priority. Because the other link connecting VLAN 2 is in Blocking mode, stations on VLAN 2 in S1 cannot communicate with stations in VLAN 2 on S2. With multiple links only one link forwards packets.

Shared servers

The Ethernet Routing Switch 2500 Series allows ports to exist in multiple VLANs for shared resources, such as servers, printers, and switch-to-switch connections. Resources can also exist in multiple VLANs on one switch, as shown in Figure 10: Multiple VLANs sharing resources on page 22.

In this example, clients on different broadcast domains share resources. The broadcasts from ports configured in VLAN 3 can be seen by all VLAN port members of VLAN 3.

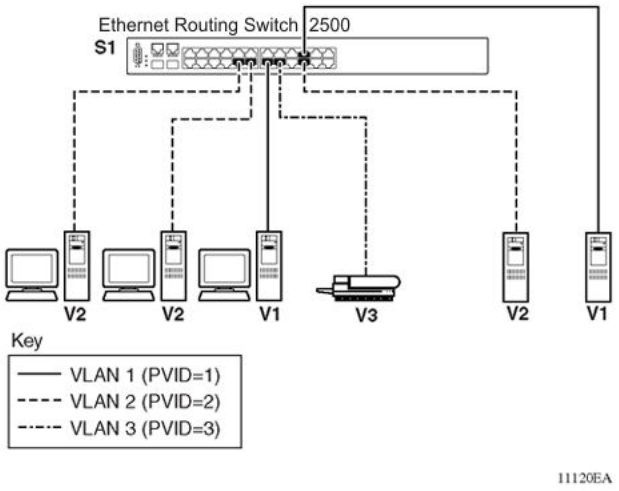

Figure 10: Multiple VLANs sharing resources

In the preceding configuration, all of the switch ports are set to participate as VLAN port members. This arrangement allows the switch to establish the appropriate broadcast domains within the switch (Figure 11: VLAN broadcast domains within the switch on page 23).

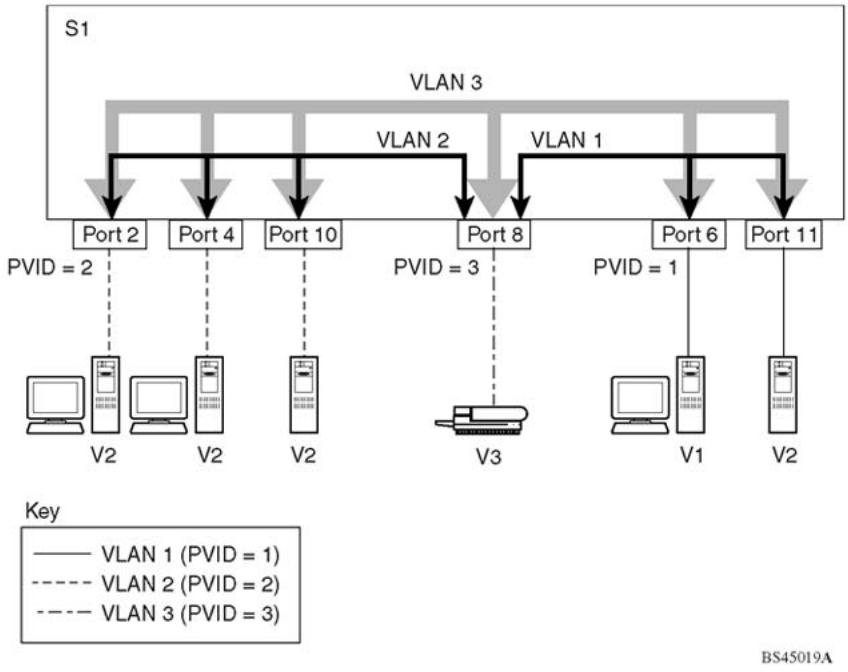

Figure 11: VLAN broadcast domains within the switch

For example, to create a broadcast domain for each VLAN shown in Figure 11: VLAN broadcast domains within the switch on page 23, configure each VLAN with a port membership, and each port with the appropriate PVID/VLAN association:

- Ports 8, 6, and 11 are untagged members of VLAN 1.
- The PVID/VLAN association for ports 6 and 11 is: PVID = 1.
- Ports 2, 4, 10, and 8 are untagged members of VLAN 2.
- The PVID/VLAN association for ports 2, 4, and 10 is: PVID = 2.
- Ports 2, 4, 10, 8, 6, and 11 are untagged members of VLAN 3.
- The PVID/VLAN association for port 8 is: PVID = 3.

VLAN workgroup summary

This section summarizes the VLAN workgroup examples discussed in the previous sections of this chapter.

As shown in Figure 12: VLAN configuration spanning multiple switches on page 24, S1 (Ethernet Routing Switch 2500 Series) is configured with multiple VLANs:

- Ports 1, 6, 11, and 12 are in VLAN 1.
- Ports 2, 3, 4, 7, and 10 are in VLAN 2.
- Port 8 is in VLAN 3.

Because S4 does not support 802.1Q tagging, a single switch port on each switch must be used for each VLAN (see VLANs spanning multiple untagged switches on page 20).

The connection to S2 requires only one link between the switches because S1 and S2 are both Ethernet Routing Switch 2500 Series devices that support 802.1Q tagging (see VLANs spanning multiple 802.1Q tagged switches on page 20).

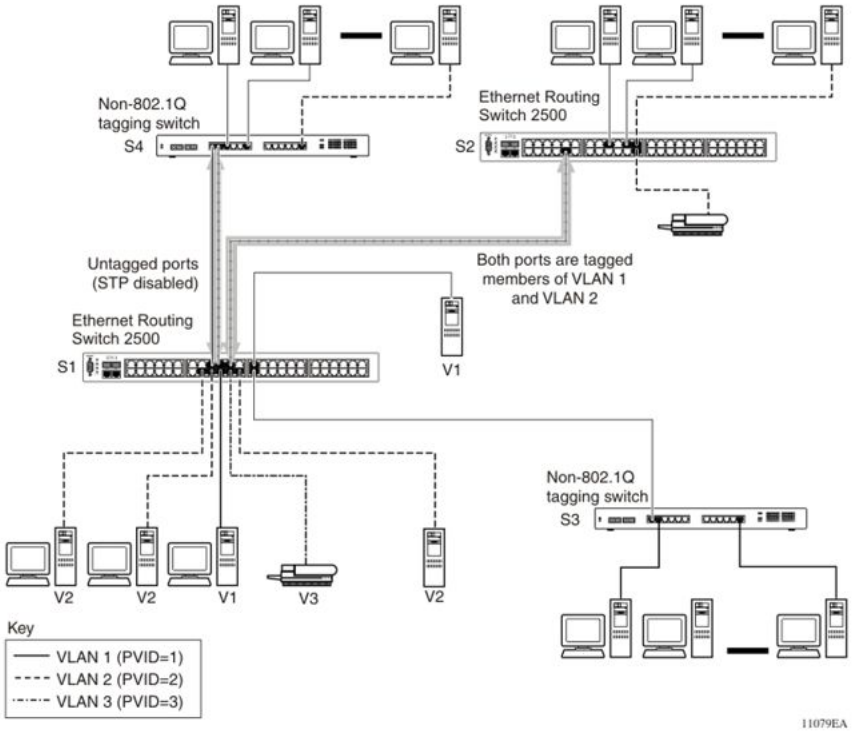

Figure 12: VLAN configuration spanning multiple switches

VLAN configuration rules

VLANs operate according to specific configuration rules. When creating VLANs, consider the following rules that determine how the configured VLAN reacts in any network topology:

- If a port is a trunk group member, all trunk members are added or deleted from the VLAN.
- All ports involved in trunking and port mirroring must have the same VLAN configuration. If a port is on a trunk with a mirroring port, the VLAN configuration cannot be changed.
- Auto PVID can be activated by creating a VLAN and enabling Auto PVID for it.

MAC Flush

You can use the MAC Flush feature to clear MAC Address entries directly from the MAC Address Table (or Forwarding Data Base). If you do not use the MAC Flush feature, you can use the following indirect methods:

- power cycling the switch
- deleting, and then recreating the VLAN
- unplugging, and then replugging the connection on the port to flush out all addresses learned on the port

MAC Flush provides the following options to flush out MAC Address entries:

- clear a single MAC Address
- clear all addresses in the MAC address table
- clear all MAC addresses from a port (or list of ports)
- clear all MAC addresses from a trunk (MLT or LAG)
- clear all MAC addresses from a particular VLAN

MAC Flush clears only dynamically learned MAC Addresses. MAC Flush does not delete MAC Addresses created by MAC Security or Port Mirroring because deletion of these MAC Addresses can affect the MAC Security or Port Mirroring function.

MAC Addresses for MAC Security or Port Mirroring have one of the following identifiers:

- AGELOCK
- SECRET
- STATIC

Higher priority tasks can delay MAC Address clearing.

Chapter 4: Spanning Tree Protocol Fundamentals

The Ethernet Routing Switch 2500 Series supports the Spanning Tree Protocol (STP) as defined in IEEE 802.1D. The Spanning Tree Protocol detects and eliminates logical loops in a bridged or switched network. When multiple paths exist, the spanning tree algorithm configures the network so that a bridge or switch uses only the most efficient path. If that path fails, the protocol automatically configures the network to make another path become active, thus sustaining network operations.

Ethernet Routing Switch 2500 Series Software Release 4.2 or later supports Rapid Spanning Tree Protocol and Multiple Spanning Tree Protocol.

This chapter contains information about the following topics:

Spanning Tree Protocol

The Ethernet Routing Switch 2500 Series supports transparent bridging by implementing the IEEE 802.1D standard. This standard is also known as the Spanning Tree Protocol (STP) and Spanning Tree Algorithm (STA) standards. STP runs on all ports to provide automatic network configuration of a loop-free topology. You can configure redundant links to provide network fault tolerance with STP.

Port states

The port will always be in one of the five states as described in the following table:

Table 2: Port states of Spanning Tree

State	Rx BPDUs	Tx BPDUs	Learn Addresses	Forward Frames
Disabled	no	no	no	no
Blocking	yes	no	no	no

State	Rx BPDUs	Tx BPDUs	Learn Addresses	Forward Frames
Listening	yes	yes	no	no
Learning	yes	yes	yes	no
Forwarding	yes	yes	yes	yes

After a switch is powered-up or reset and the initialization process is completed, all the ports are transformed from the Disabled state to the Blocking state.

If a port is not connected, the port remains in the Forwarding state until it is connected.

If you connect a station to a port, the port does not forward packets immediately. You must wait for the port to transit through the Listening and Learning states to have access to any resources located on another segment.

If you connect a hub or another bridging device to a port, it creates a loop in the network topology and a broadcast storm can occur. This problem can occur if one of the ports causing the loop is in the Forwarding state instead of the Blocking state. The loop will disappear when this port receives a BPDU frame from a higher priority port.

Use the MIB variable dot1dStpPortEnable to disable or enable a port. A port is enabled by default. In this mode of operation, the port is in one of the following STP states:

- Blocking
- Listening
- Learning
- Forwarding

If you disable a port, it will not forward any frames and will not participate in the Spanning Tree Algorithm and Spanning Tree Protocol.

STP port mode

With the STP port mode feature, a switch port can maintain participation in an STP if the port is moved from one VLAN to another.

When the STP port mode is configured to auto and a port which does not belong to any VLAN is added to a VLAN, the STP participation of the port is automatically enabled. If the STP port mode is configured to normal and a port which does not belong to any VLAN is added to a VLAN, the STP participation of the port is disabled.

STP 802.1d compliance mode

STP 802.1d compliance mode can ensure that STP conforms to the IEEE 802.1d standard. IEEE 802.1d indicates that when a port link fails, the STP state of the port should stay in

Forwarding mode. When STP 802.1d compliance mode is enabled, the switch is provided a fast recovery mechanism for a port that frequently changes state from up to down.

STP 802.1d compliance mode ensures that STP conforms to the IEEE 802.1d standard. When the STP 802.1d compliance mode is disabled, the switch is provided a fast recovery mechanism for a port that frequently changes state from up to down. This fast recovery mechanism does not comply with the IEEE 802.1d standard, so when STP 802.1d compliance mode is enabled, the fast recovery mechanism is no longer available and the passing from blocking to forwarding state is done through listening and learning states. When a port link fails, the STP state of the port is Forwarding if STP 802.1d compliance mode is disabled and the STP state of the port is Disabled if STP 802.1d compliance mode is enabled.

Aging of dynamic entries in Forwarding Database

Dynamic MAC address entries are automatically removed from the Forwarding Database after a specified time.

If the network topology did not change, the aging timeout value is specified by the dot1dTpAgingTime MIB variable. This can be configured through the user interface console. The range of applicable values specified in the IEEE standard is 10 to 1000000 seconds, whereas Avaya recommends a default value 300 seconds.

If the root bridge notifies other bringing devices of topology changes, to other bridging devices, a short aging timeout value is used. The timeout value is set equal to the Forward Delay parameter contained in BPDUs originating from the root. The range of values for the Forward Delay parameter specified in the IEEE standard is 4 to 30 seconds. Avaya recommend a default value is of 15 seconds.

Port path cost

You can assign the path cost or the switch can automatically calculate the path cost associated with a port. By default the path cost is automatically calculated. Also by default, the cost of a given link is originally specified (IEEE90) to be inversely proportional to the data rate of the link. Thus, a 10 Mb/s Ethernet has a link cost of 100. This formula does not work well for Gigabit Ethernet or even for emerging technologies such as packets-over-SONET at OC-48 rates and above.

Table 3: Path Cost Values on page 29 describes a range of values for a given data rate, and a recommended value that has a nonlinear relationship between link cost and data rate for very high-speed LANs.

Table 3: Path Cost Values

Data rate	Recommended link cost range	Recommended link cost value
10 Mb/s	50 to 600	100

Data rate	Recommended link cost range	Recommended link cost value
100 Mb/s	10 to 60	10
1 Gb/s	3 to 10	1
10 Gb/s	1 to 5	1

The valid range for path cost values is between 0 and 65535. If you enter a value between 1 and 65535, the port path cost is set to the new value.

802.1t path cost calculation

In release 4.0 software and later, you can set the switch to calculate the STG path cost using either the IEEE 802.1d standard or the IEEE 802.1t standard. The 802.1t standard is a maintenance extension to the 802.1d standard.

Rapid Spanning Tree Protocol

The current Spanning Tree implementation in Ethernet Routing Switch 2500 Series is based on IEEE 802.1d, which is slow to respond to a topology change in the network (such as a dysfunctional link in a network). The Rapid Spanning Tree Protocol (RSTP or IEEE 802.1w) reduces the recovery time after a network breakdown. In certain configurations the RSPT recovery time is less than 1 second. It also maintains a backward compatibility with the IEEE 802.1d, which was the Spanning Tree implementation prior to RSTP. The backward compatibility can be maintained by configuring a port to be in STP compatible mode. A port operating in the STP compatible mode transmits and receives only STP BPDUs and drops any RSTP BPDUs.

RSTP also reduces the amount of flooding in the network by enhancing the way Topology Change Notification (TCN) packet is generated.

Multiple Spanning Tree Protocol

With Multiple Spanning Tree Protocol (MSTP or IEEE 802.1s), you can configure multiple instances of RSTP on the same switch. Each RSTP instance can include one or more VLANs. The operation of the MSTP is similar to the current Avaya proprietary MSTP.

The Ethernet Routing Switch 2500 Series use RSTP and MSTP to achieve the following:

- Reduce converging time from 30 seconds to less than 1 or 2 seconds when there is topology change in the network (such as, a port in or out of service).
- Eliminate unnecessary flushing of the MAC database and flooding of traffic to the network, using new Topology Change mechanism.
- Backward compatibility with other switches that run legacy 802.1d STP.
- Under MSTP mode, eight instances of RSTP can be supported simultaneously. Instance 0 or CIST is the default group, which includes default VLAN 1. Instances 1 to 7 are called MSTIs 1 to 7.
- You can configure the switch to run avayaStpg, RSTP, or MSTP configuration.

Interoperability with legacy STP

RSTP provides a new parameter—Force Version for backward compatibility with legacy STP. You can configure a port in either STP compatible mode or RSTP mode.

- An STP compatible port transmits and receives only STP BPDUs. Any RSTP BPDU that the port receives in this mode will be discarded.
- An RSTP compatible port transmits and receives only RSTP BPDU. If an RSTP port receives a STP BPDU it becomes an STP port. User intervention is required to bring this port back to RSTP mode. This process is called Port Protocol Migration.

Differences in port roles

RSTP is an enhanced version of STP. These two protocols have almost the same set of parameters.

Table 4: Differences in port roles for STP and RSTP on page 31 lists the differences in port roles for STP and RSTP. STP supports two port roles while RSTP supports four port roles.

Table 4: Differences in port roles for STP and RSTP

Port role	STP	RSTP	Description
Root	Yes	Yes	This port is receiving a better BPDU than its own and it has the best path to reach the Root. Root port is in Forwarding state.
Designated	Yes	Yes	This port has the best BPDU on the segment. Designated port is in Forwarding state.

Port role	STP	RSTP	Description
Alternate	No	Yes	This port is receiving a better BPDU than its own BPDU and there is a Root port within the same switch. Alternate port is in Discarding state.
Backup	No	Yes	This port is receiving a better BPDU than its own BPDU and this BPDU is from another port within the same switch. Backup port is in Discarding state.

Edge port

Edge port is a new parameter that RSTP supports. When you connect a port to a nonswitch device such as a PC or a workstation, you must configure it as an Edge port. An active Edge port goes directly to Forwarding state without any delay. An Edge port becomes a non-Edge port if it receives a BPDU.

Path cost values

RSTP and MSTP recommend new path cost values that support a wide range of link speeds. Table 5: Recommended path cost values on page 32 lists the recommended path cost values.

Table 5: Recommended path cost values

Link speed	Recommended value
Less than or equal 100Kb/s 1 Mb/s 10 Mb/s 100 Mb/s	200 000 000 20 000 000 2 000 000 200 000
1 Gb/s 10 Gb/s 100 Gb/s	20 000 2 000 200
1 Tb/s 10 Tb/s	20 2

Rapid convergent

In RSTP and MSTP the environment root port or the designated port can ask its peer for permission to go to the Forwarding state. If the peer agrees then the root port can move to the Forwarding state without any delay. This procedure is called negotiation process.

RSTP and MSTP also lets the switch send information received on a port immediately if the port becomes dysfunctional instead of waiting for the Maximum Age time.

The following example illustrates how an RSTP port moves rapidly to Forwarding state without the risk of creating a loop in the network.

Switch A: ports 1 and 2 are in full duplex. Port 2 is an Edge port

Switch B: ports 1, 2 and 3 are in full duplex. Port 2 is an Edge port.

Switch C: ports 1 and 2 are in full duplex. Port 2 is an Edge port.

Switch A is the Root.

Negotiation process

After power up, all ports assume the role as Designated ports. All ports are in the Discarding state except Edge ports. Edge ports go directly to Forwarding state without delay.

Switch A port 1 and switch B port 1 exchange BPDUs. Switch A is the Root and switch A port 1 is the Designated port. Switch B learns that switch A has better priority. Switch B port 1 becomes Root port. Both switch A port 1 and switch B port 1 are still in Discarding state.

Switch A starts negotiation process by sending BPDU with proposal bit set.

Switch B receives the proposal BPDU and sets its non-Edge ports to Discarding state. This operation is called the synchronization process.

Switch B sends a BPDU with the agreement bit set to switch A.

Switch A sets port 1 to Forwarding state and switch B sets port 1 to Forwarding state. PC 1 and PC 2 communicate with each other.

The negotiation process now moves down to switch B port 3 and its partner port.

PC 3 cannot communicate with either PC 1 or PC 2 until the negotiation process between switch B and switch C is complete.

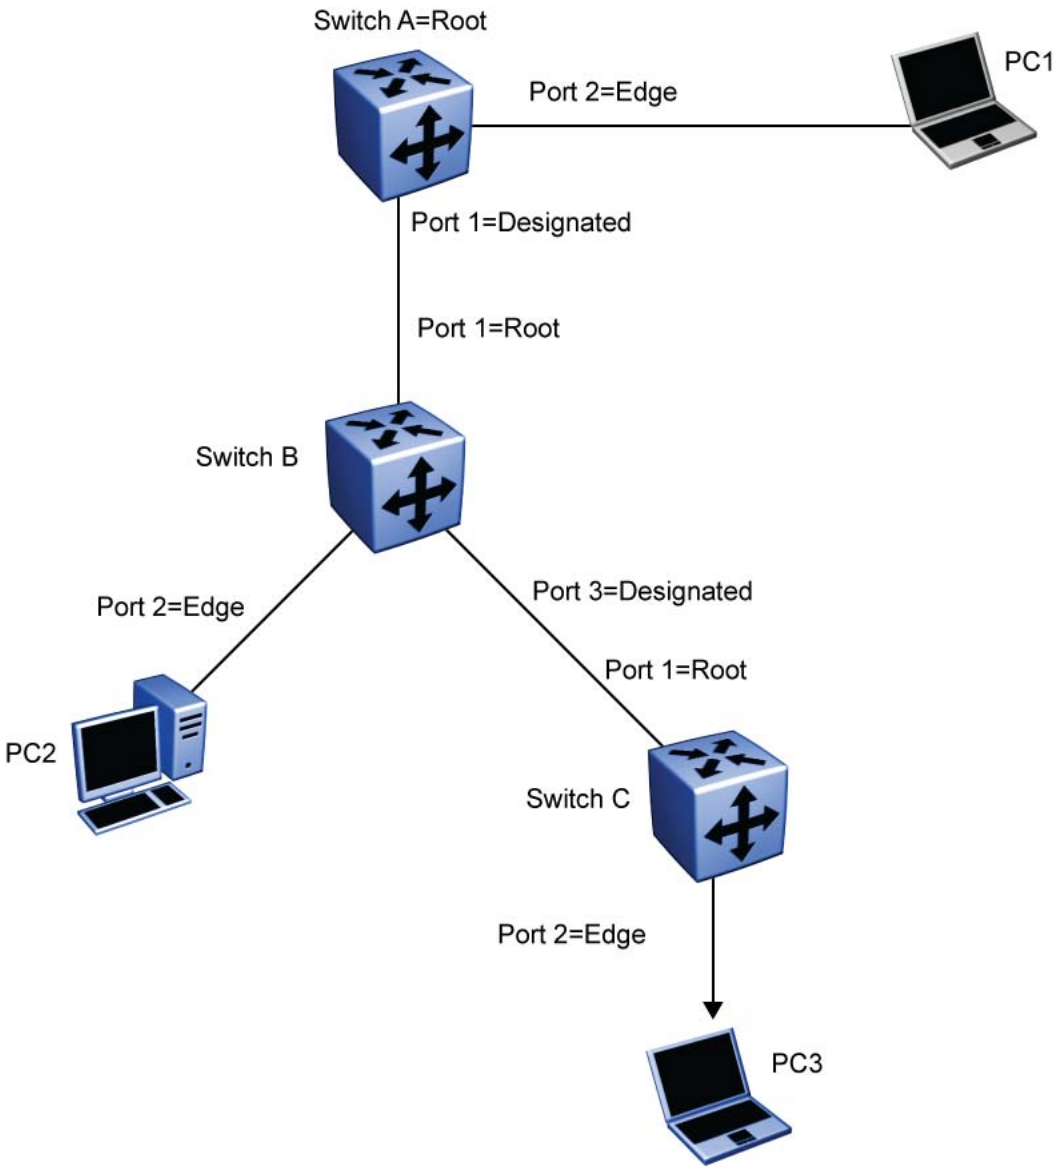

Figure 13: Negotiation process

Spanning Tree BPDU Filtering

Release 4.2 or later Software supports the BPDU-Filtering feature for STPG, RSTP, and MSTP.

The Spanning Tree Protocol detects and eliminates logical loops in a bridged or switched network. Any bridge that participates in the spanning tree exchanges information with other bridges using configuration messages known as Bridge Protocol Data Units (BPDU). Based

on the BPDU information exchange, the bridge with the lowest bridge ID becomes the root. This process is called the root selection process.

Typically, when a new bridge joins the spanning tree or an existing bridge leaves the spanning tree, the root selection process is repeated and a new root is selected.

The BPDU-Filtering feature allows the network administrator to achieve the following:

- Block an unwanted root selection process when an edge device, such as a laptop running Linux and enabled with STP, is added to the network. This prevents unknown devices from influencing an existing spanning tree topology.
- Block the flooding of BPDUs from an unknown device.

Note:

The STP BPDU-Filtering feature is not supported on Multi-Link Trunk (MLT) ports.

When a port has BPDU-Filtering enabled and it receives an STP BPDU, the following actions take place:

- The port is immediately put in the operational disabled state.
- A trap is generated and the following log message is written to the log: `BPDU received on port with BPDU-Filtering enabled. Port <x> has been disabled`
- The port timer starts.
- The port stays in the operational disabled state until the port timer expires.

If the timer is disabled or the switch is reset before the timer expires, the port remains in the disabled state. Similarly, if a user disables BPDU-Filtering while the timer is running, the timer is stopped and that port stays in the disabled state. In this case, you must then manually enable the port to bring it back to the normal mode.

You can enable and disable the BPDU-Filtering feature on a per-port basis. The BPDU-Filtering timer is user-configurable for each port and has a valid range of between 10 and 65535 seconds. The port timer is disabled if it is configured as 0.

For details about configuring BPDU Filtering, refer to: Configuring STP BPDU Filtering using EDM on page 159 and Configuring STP BPDU filtering on page 91.

Chapter 5: Multi-Link Trunking Fundamentals

This chapter contains information about the following topics:

About Multi-Link Trunking

The Multi-Link Trunking (MLT) feature is a point to point link aggregation function that allows you to group multiple switch ports together, when forming a link to another switch or server. This provides additional link redundancy and increases the aggregate throughput of the interconnection between two devices.

The Ethernet Routing Switch 2500 Series can be configured with up to six (6) Multi-Link Trunk groups, of up to four (4) links within each group. Multi-Link Trunking software detects broken trunk links and redirects traffic from the broken trunk link(s) to other trunk members within that trunk.

The MLT feature supports the grouping of ports on one switch or across multiple switches in a switch stack. This provides additional link redundancy while also building a higher bandwidth connection between two network devices, with the traffic load balanced across the physical ports in the trunk group.

Trunking can be described in the following terms:

- Network Trunk (NT) - A NT is connected to another internetworking device.
- Server Trunk (ST) - A ST is attached to a server that utilizes the same MAC address on each of its links.

The two basic switching requirements of MLTs are:

- The ability to treat multiple links as a single one for the purposes of learning and migration.
- The ability to select one of the member paths as the destination for a forwarding function without sending any duplicate packets.

MLT operation

Ethernet Routing Switch 2500 Series supports a maximum of six trunks, scaling up to four ports per trunk. The MLT operation is based on the concept of trunk groups. A trunk group is

a collection of ports that represent a single link for learning, forwarding and other bridge functions.

Forwarding model

The trunk forwarding function is based on the following:

- Destination Address (DA)
- Source Address (SA)

The forwarding model has two modes, Basic and Advanced. To select the egress link in a trunk configuration, Basic mode uses the source and destination MAC addresses of learned packets, while Advanced mode uses the source and destination IP addresses.

A maximum of four ports will be assigned to a trunk. The source address is associated with the trunk group rather than the individual port it was learned on. From here, the forwarding function points the packets to that particular trunk group.

For proper network operation, packets cannot be replicated to more than one port of a trunk group. The operation that creates this selection is based on the SA. SA selects one of the possible egress ports that is a member of the trunk group. For any DA, the egress path will always be defined by the SA.

Packets to a certain DA can appear on any member link of the trunk. Packets with the same SA always appear on the same egress port irrespective of DA. The exception to this is when the BCAST/MCAST/DLF traffic is sent out using the same port within the MLT regardless of the SA.

MLT configuration examples

You can use the Trunk Configuration screen to create switch-to-switch and switch-to-server Multi-Link Trunk links (Figure 14: Switch-to-switch trunk configuration example on page 39 and Figure 15: Switch-to-server trunk configuration example on page 40).

Figure 14: Switch-to-switch trunk configuration example on page 39 shows two trunks (T1 and T2) connecting Switch S1 to switches S2 and S3.

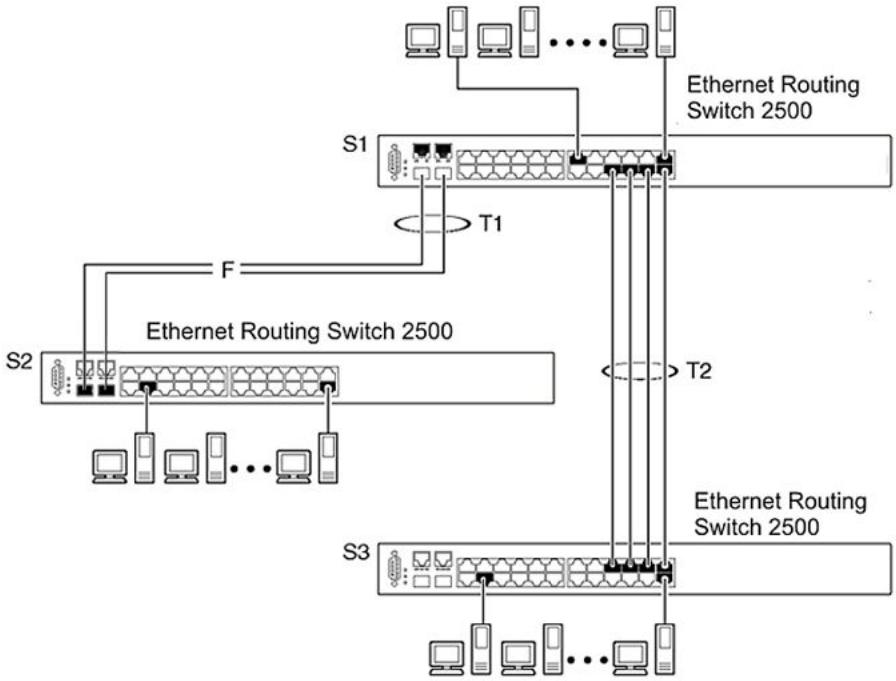

Figure 14: Switch-to-switch trunk configuration example

You can configure each trunk shown in Figure 14: Switch-to-switch trunk configuration example on page 39 with a maximum of four ports on the Ethernet Routing Switch 2500 Series to provide 400 Mb/s aggregate bandwidth through T2 or 2Gb/s aggregate bandwidth through T1, in full-duplex mode. As shown in the example, creating a Multi-Link Trunk can supply additional bandwidth required to improve the performance when the traffic between switch-to-switch connections approach single port bandwidth limitations.

Figure 15: Switch-to-server trunk configuration example on page 40 shows a typical switch-to-server trunk configuration. In this example, file server FS1 uses dual MAC addresses, using one MAC address for each network interface card (NIC). For this reason, FS1 does not require a trunk assignment. FS2 is a single MAC server (with a four-port NIC) and is set up as trunk configuration T1.

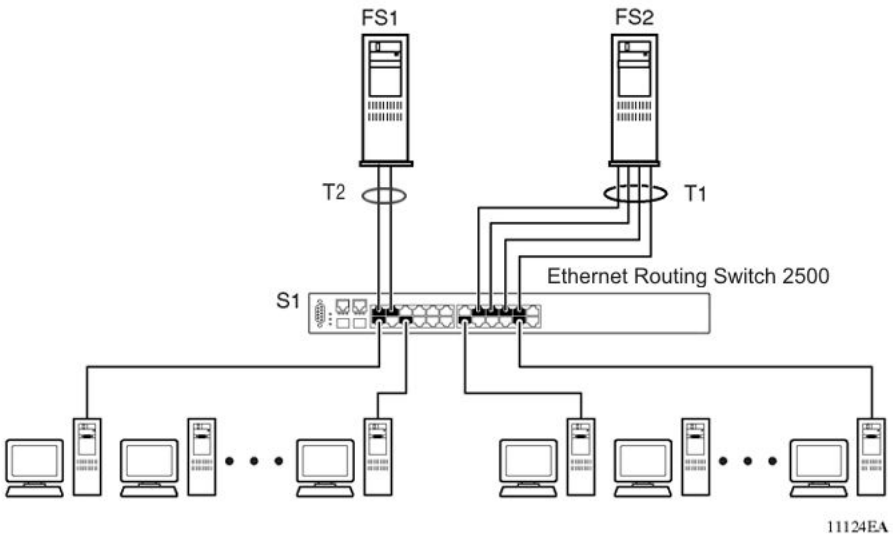

Figure 15: Switch-to-server trunk configuration example

Client/server configuration using Multi-Link Trunks

Figure 16: Client/server configuration example on page 41 shows an example of how Multi-Link Trunking can be used in a client/server configuration. In this example, both servers connect directly to Switch S1. FS2 is connected through a trunk configuration (T1). The switch-to-switch connections are through trunks (T3, T4, and T5).

Clients accessing data from the servers (FS1 and FS2) are provided with maximized bandwidth through trunks T1, T2, T3, T4, and T5. Trunk members (the ports making up each trunk) do not have to be consecutive switch ports; you can select ports randomly, as shown by T5.

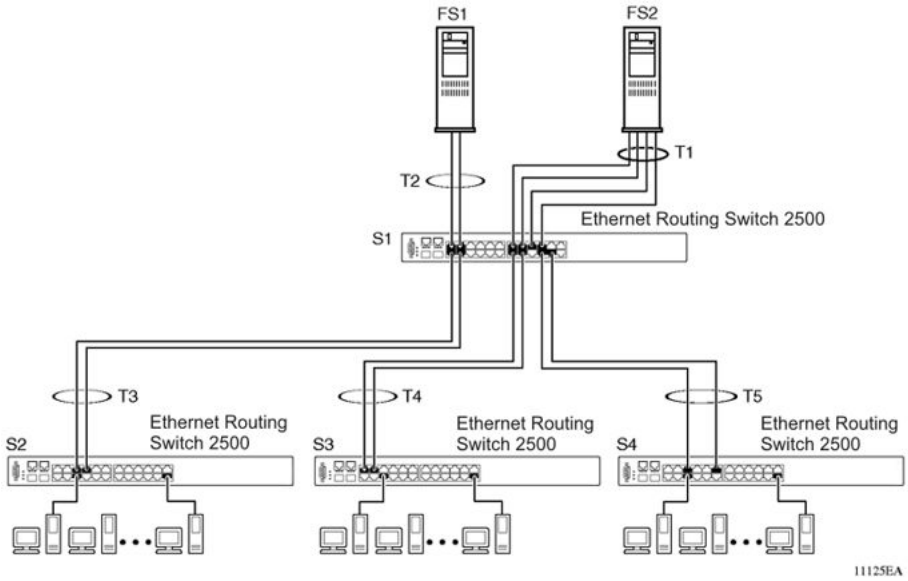

Figure 16: Client/server configuration example

For detailed information about configuring trunks, see Multi-Link Trunking configuration using ACLI on page 93 and Configuring Multi-Link Trunking using Enterprise Device Manager on page 189.

Before you configure trunks

When you create and enable a trunk, the trunk members (switch ports) take on certain settings necessary for correct operation of the Multi-Link Trunking feature.

Before you configure your Multi-Link Trunk, you must consider these settings, along with specific configuration rules, as follows:

1. Read the configuration rules provided in the next section, Spanning tree considerations for Multi-Link Trunks on page 42.
2. Determine which switch ports (up to four) are to become trunk members (the specific ports making up the trunk). A minimum of two ports are required for each trunk.

   Disabled ports can belong to MLTs. To enable traffic to flow to your configured MLT ports, ensure that the chosen switch ports are set to Enabled.

   Trunk member ports must have the same VLAN and VLACP configuration. LACP should not be enabled on the selected trunk ports.
3. All network cabling should be complete and stable before configuring any trunks, to avoid configuration errors.

4. Consider how the existing spanning tree reacts to the new trunk configuration (see Spanning tree considerations for Multi-Link Trunks on page 42).
5. Consider how existing VLANs are affected by the addition of a trunk.

Spanning tree considerations for Multi-Link Trunks

The spanning tree Path Cost parameter is recalculated based on the aggregate bandwidth of the trunk. For example, Figure 17: Path Cost arbitration example on page 42 shows a four-port trunk (T1) with two port members operating at 100 Mb/s and two at 10 Mb/s. Trunk T1 provides an aggregate bandwidth of 220 Mb/s. The Path Cost for T1 is 4 (Path Cost = 1000/ LAN speed, in Mb/s). Another three-port trunk (T2) is configured with an aggregate bandwidth of 210 Mb/s, with a comparable Path Cost of 4. When the path cost calculation for both trunks is equal, the spanning tree software chooses the trunk with the lowest Spanning Tree PortID, regardless of the aggregate bandwidth.

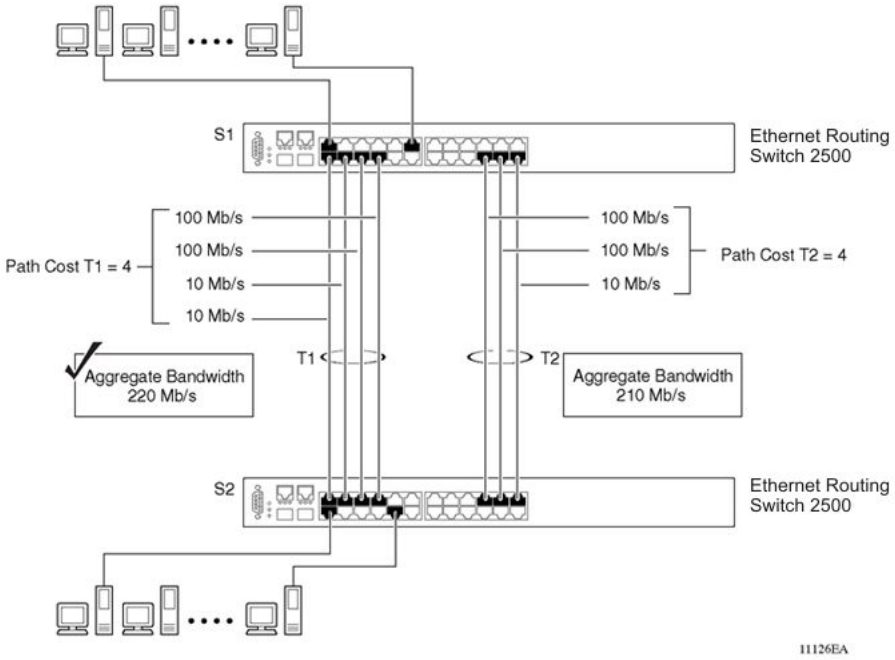

Figure 17: Path Cost arbitration example

Additional tips about the Multi-Link Trunking feature

When you create a Multi-Link Trunk, the individual trunk members (the specific ports that make up the trunk) logically connect and react as a single entity. For example, if you change spanning tree parameters for any trunk member, the spanning tree parameters for all trunk members change.

The trunk is viewed by management stations as a single spanning tree port. The spanning tree port is represented by the trunk member with the lowest port number. For example, if ports 13,

14, 15, and 16 are trunk members of trunk T1, the management station views trunk T1 as spanning tree port 13.

Important:

At boot time, the agent verifies the setting consistency for various applications (like Rate Limiting, EAP, and Port Mirroring) on the MLT ports. MLT is disabled if they are inconsistent.

See also Quick configuration for Multi-Link Trunking on page 223 for a configuration flowchart that can help you use this feature.

MLT enable or disable whole trunk

The MLT enable or disable whole trunk feature is user configurable and can be enabled or disabled switch-wide with a single CLI command. The feature is disabled by default. With the MLT whole trunk disabled, you can enable or disable MLT or DMLT groups, and the operational states of the bundled links do not change. In this configuration, a network traffic loop can occur when you disable MLT or DMLT groups. The switch supports the ability to change this operational mode using the MLT whole trunk feature.

If you enable the MLT whole trunk feature, the underlying state of the port changes to reflect the state of the MLT or DMLT bundle regardless of the previous status. With the MLT whole trunk enabled, you can disable the MLT or DMLT and all links that are part of the MLT group are disabled except for the Default Forwarding Link (DLF), which remains active to prevent loss of connectivity to the switch or stack. The DLF link is typically the lowest numbered port of an active MLT or DMLT link. Conversly, if you enable the MLT or DMLT, all links will become active.

You can enable or disable individual links of a MLT or DMLT if the MLT whole trunk feature is enabled.

Important:

For network configuration, Avaya recommends that you enable the MLT whole trunk feature.

Chapter 6: LACP and VLACP Fundamentals

IEEE 802.3ad Link Aggregation

You can create and manage a trunk group with Link Aggregation (LA) . You can control and configure a trunk group automatically using the Link Aggregation Control Protocol (LACP).

The LACP, defined by the IEEE 802.3ad standard, allows the switch to learn the presence and capabilities of a remote switch by exchanging information with the remote switch before a trunk group is formed. Either switch can accept or reject the aggregation request with the far end on a per port basis. A link that can not join a trunk group operates as an individual link. 802.3ad provides an industry standard method for bundling multiple links together to form a single trunk between two networking devices. Trunks that conform to the 802.3ad standard are Link Aggregation Groups (LAGs). Release 4.2 or later software supports 2 types of trunks:

- Dynamic LAG
- MLT

A trunk group that is formed by Link Aggregation is called a Link Aggregation group (LAG), and a trunk group that is formed by Ethernet Multi-link Trunking is called a Multi-link trunk (MLT) group.

The Ethernet Routing Switch 2500 Series supports both Link Aggregation groups and Multi-link trunks. By default, Link Aggregation is set to disabled on all ports. A Link Aggregation group or trunk group can be created or deleted automatically using Link Aggregation Control Protocol (LACP).

The maximum number of Link Aggregation and MLT groups is six, and the maximum number of active links per group is four. Link Aggregation allows more than four links to be configured in one Link Aggregation group (LAG).

The first four high priority links are active links and together they form a trunk group. The remaining low priority links remain in standby mode. When one of the active links goes down, one of the standby links becomes active and is added to the trunk group.

The failover process is as follows:

- The down link is removed from the trunk group
- The highest priority standby link is added to the trunk group

Important:

The STP participation for an active MLT or LAG trunk always overrides the STP participation previously configured for individual ports. If a user changes the STP participation on

individual trunk ports after the trunk is disabled, the port STP participation will be overridden by the Trunk's STP participation after the trunk is enabled again.

There can be a temporary delay in traffic flow due to the switching of links. If the active link goes down and there is no standby link, the traffic is re-routed to the remaining active links with a minimal delay in time.

Half duplex links are not allowed in LAG, and all links in a LAG must have the same speed.

802.3 Link Aggregation is available through the Avaya Command Line Interface (ACLI). The ACLI supports the following commands:

The following ACLI commands can be executed to enable, disable, or set default values for LACP on a port:

- **lacp aggregation [port <portlist>] enable**
- **no lacp aggregation [port <portlist>] enable**
- **default lacp aggregation [port <portlist>] enable**

To specify the LACP mode:

- **lacp mode [port <portlist>] {off | passive | active}**
- **default lacp mode [port <portlist>]**

To assign an administrative key value to a port:

lacp key [port <portlist>] <1-4095>

To specify the port priority:

- **lacp priority [port <portlist>] <0-255>**
- **default lacp priority [port <portlist>]**

To set port time-out:

- **lacp time-out-time [port <portlist>] {short | long}**
- **default lacp time-out-time [port <portlist>]**

To set LACP system priority:

- **lacp system-priority [0-65535]**
- **default lacp system-priority**

ACLI Show commands for LACP:

- **show lacp aggr**
- **show lacp port[<portList>]**
- **show lacp port aggregator**
- **show lacp debug member [portlist]**
- **show lacp system**

- **show lacp stats [port <portlist>]**
- **show lacp stats aggregator**
- **lacp clear-stats**

For more information about the syntax and parameters of the ACLI commands, see Configuring Link Aggregation Group on page 96.

VLACP

Many enterprise networks require that trunk links provide subsecond failover to the redundant link after a failure occurs at the local or remote endpoint. This requirement can be met after both ends of the link are informed of any loss of communication.

Virtual Link Aggregation Control Protocol (VLACP), an LACP extension, is a Layer 2 handshaking protocol that provides end-to-end failure detection between two physical Ethernet interfaces. It allows the switch to detect unidirectional or bidirectional link failures.

Virtual LACP (VLACP) overview

While Ethernet has been extended to detect remote link failures through functions such as Remote Fault Indication and Far End Fault Indication mechanisms, a limitation of these functions is that they terminate at the next Ethernet hop. Therefore, failures cannot be determined on an end-to-end basis.

Enterprise networks can connect their aggregated Ethernet trunk groups through a service provider network connection (for example, through a VPN), but far-end failures cannot be signaled with Ethernet-based functions that operate end-to-end through a service provider cloud.

In the following example, the MLT (between Enterprise switches S1 and S2) extends through the service provider (SP) network.

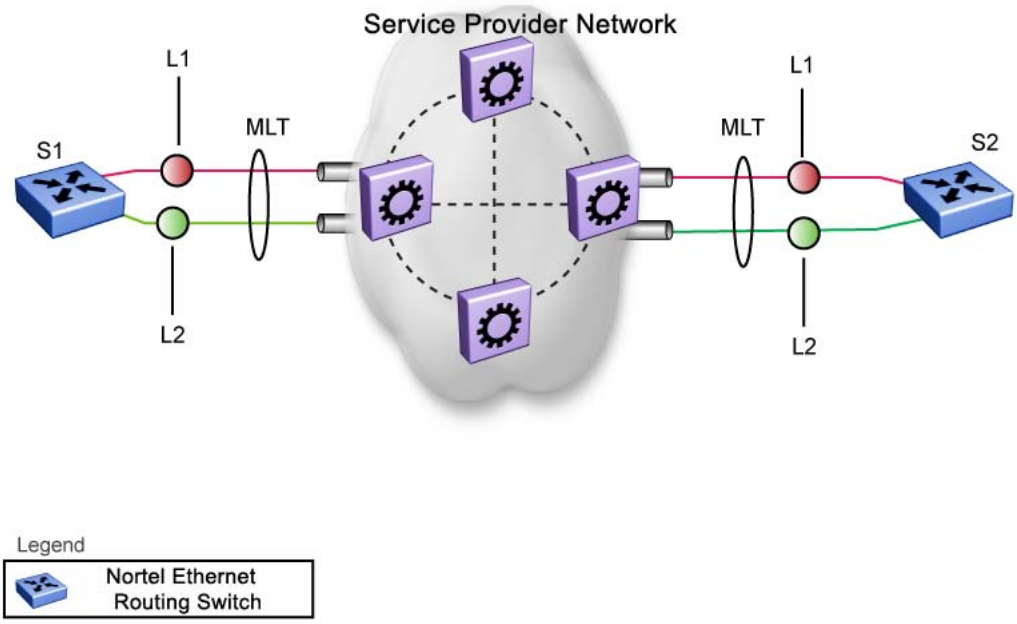

As shown in the next example, if the L2 link on S1 (S1/L2) fails, the link-down failure is not propagated over the SP network to S2. Thus, S2 continues to send traffic over the S2/L2 link, which is black-holed because the S1/L2 link has failed.

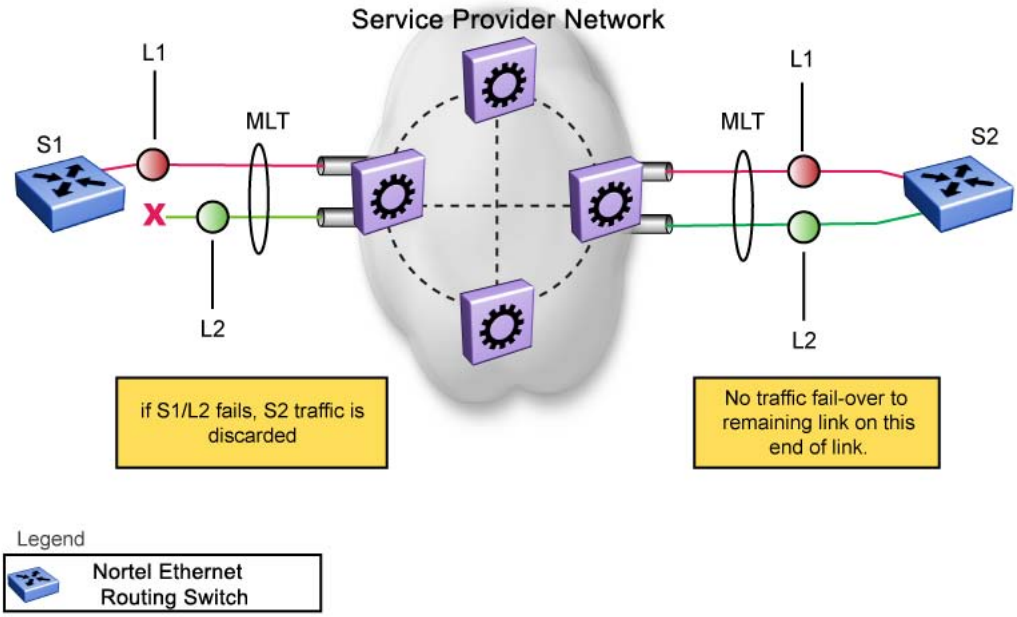

Note that LACP, as defined by IEEE, is a protocol that exists between two bridge endpoints; therefore, the LACPDUs are terminated at the next (SP) interface.

Avaya has developed an extension to LACP, which is called Virtual LACP (VLACP). This extension can provide an end-to-end failure detection mechanism. With VLACP, far-end failures can be detected allowing an MLT to fail over properly when end-to-end connectivity is not guaranteed for certain links in an aggregation group.

VLACP features

This section provides a summary of some of the key features of VLACP:

- VLACP is configured per port. A port can be an individual port or a member of an MLT.
- When you set VLACP parameters for a trunk port, the settings are applied to all trunk members.
- For VLACP to operate properly, there must be a logical point-to-point connection (Layer 2 tunnel) between the two endpoints.
- VLACP does not work for point-to-multipoint connections.
- On each port that has VLACP enabled, VLACPDUs are sent periodically. If VLACPDUs are not received on a particular link, that link is taken down after a configurable timeout period.
- For the current software release, VLACP is supported on Ethernet interfaces only.
- VLACP can run independently as a port-to-port protocol or on top of MLT or LACP protocol.
- VLACP packets are untagged because they operate at the port level and not the VLAN level.
- The Destination Mac Address used in VLACPDUs is configurable. The MAC Address must be a multicast MAC Address so that it is always flooded. This allows the exchange of VLACPDUs from end to end.

Troubleshooting

Error logs are created for the following failures and errors:

- An incorrect PDU, such as wrong destination MAC addresses received
- An inability to enable VLACP on a port due to unallowable Destination MAC addresses
- A port index that is out of range
- A port was blocked by VLACP (a log message is also generated after the port is unblocked)

Chapter 7: ADAC Fundamentals

Ethernet Routing Switch 2500 Series supports the Auto-Detection and Auto-Configuration (ADAC) of Avaya IP Phones. With ADAC, you can automatically configure the switch to support and prioritize IP Phone traffic.

When ADAC is enabled and a Avaya IP Phone is connected to the switch, the switch automatically configures the VLAN, port, and Quality of Service (QoS) settings necessary for the transmission of signal and voice between the Avaya IP Phone and the switch.

ADAC can configure the switch whether the switch is directly connected to the Call Server (through the Call Server port) or is indirectly connected to the Call Server using a network uplink (through the Uplink port).

ADAC has three separate operating modes to meet the requirements of different networks:

- **Untagged-Frames-Basic:**

  Use this mode when you want a basic configuration only and the IP Phones are sending untagged traffic.

- **Untagged-Frames-Advanced:**

  Use this mode when you want an advanced configuration and the IP Phones are sending untagged traffic. In this mode, ADAC creates a Voice VLAN that includes the Call Server or Uplink port, as applicable, and all telephony ports. All tagging, PVID settings, and traffic prioritization are configured automatically.

- **Tagged Frames:**

  Use this mode when you want an advanced configuration and the IP Phones are sending tagged traffic. You can also use tagged frames to support devices other than IP Phones. This mode provides the same configuration as the Untagged-Frames-Advanced mode, but with tagged frames. As with the Untagged-Frames-Advanced mode, ADAC creates a Voice VLAN that includes the Call Server or Uplink port, as applicable, and all telephony ports. All tagging, PVID settings, and traffic prioritization are configured automatically.

ADAC operation

The following sections provide detailed explanations of ADAC operation.

Auto-Detection of Avaya IP Phones

When a Avaya IP Phone is connected to a switch and is powered on, the switch automatically detects the IP Phone, and then begins the auto-configuration of the IP Phone. An ADAC lookup is also performed each time a MAC address is learned, migrated, or aged-out and removed.

When you enable auto-detection on a port, the port also becomes operationally enabled. Similarly, after you disable auto-detection on a port, the port is operationally disabled. A port can also be operationally disabled if the port maximum of 32 devices is reached. If the port limit is reached, a trap will be sent (if ADAC traps are enabled) and auto-configuration will also be removed. To put the port back into the operational state, disable and then re-enable auto-detection on the affected port. ADAC supports a maximum of 32 devices (both IP phones and non-phones) per port.

There are two ways to use ADAC to automatically detect IP Phones. You can enable one or the other or both of these methods on a port-by-port basis, as long as at least one detection mechanism remains enabled.

The detection mechanism can be selected either before enabling auto-detection on the port, or if ADAC is globally disabled

The two methods of auto-detection are by MAC address or using LLDP (IEEE 802.1AB).

Auto-detection by MAC address is based on using predefined MAC addresses to determine that the specified port is connected to a Avaya IP phone. For more information and the list of defined MAC address ranges, see Auto-Detection by MAC address on page 52

Auto-detection by LLDP allows the system to detect IP phones with MAC addresses outside the list of default MAC address ranges as long as they can be identified as an IP phone by LLDP, regardless of their MAC addresses. For more information about auto-detection by LLDP, see Auto-Detection by LLDP (IEEE 802.1AB) on page 54.

You can enable either of these detection mechanisms or both on each individual port. At least one of these detection methods must be enabled on each port.

Auto-Detection by MAC address

When this feature is enabled on a port, the switch checks all MAC addresses of packets received on the port. If a received MAC address falls within the range of known Avaya IP Phone MAC addresses, ADAC determines that the specified port is connected to a Avaya IP Phone and initiates the required configuration. ADAC is supported for a maximum of 32 devices per port, but in most cases, there will be only one IP phone and one PC on each port.

Table 6: Default ADAC MAC address ranges on page 53shows a list of the default MAC address ranges.

Table 6: Default ADAC MAC address ranges

Lower End	Higher End
00-0A-E4-01-10-20	00-0A-E4-01-23-A7
00-0A-E4-01-70-EC	00-0A-E4-01-84-73
00-0A-E4-01-A1-C8	00-0A-E4-01-AD-7F
00-0A-E4-01-DA-4E	00-0A-E4-01-ED-D5
00-0A-E4-02-1E-D4	00-0A-E4-02-32-5B
00-0A-E4-02-5D-22	00-0A-E4-02-70-A9
00-0A-E4-02-D8-AE	00-0A-E4-02-FF-BD
00-0A-E4-03-87-E4	00-0A-E4-03-89-0F
00-0A-E4-03-90-E0	00-0A-E4-03-B7-EF
00-0A-E4-04-1A-56	00-0A-E4-04-41-65
00-0A-E4-04-80-E8	00-0A-E4-04-A7-F7
00-0A-E4-04-D2-FC	00-0A-E4-05-48-2B
00-0A-E4-05-B7-DF	00-0A-E4-06-05-FE
00-0A-E4-06-55-EC	00-0A-E4-07-19-3B
00-0A-E4-08-0A-02	00-0A-E4-08-7F-31
00-0A-E4-08-B2-89	00-0A-E4-09-75-D8
00-0A-E4-09-BB-9D	00-0A-E4-09-CF-24
00-0A-E4-09-FC-2B	00-0A-E4-0A-71-5A
00-0A-E4-0A-9D-DA	00-0A-E4-0B-61-29
00-0A-E4-0B-BB-FC	00-0A-E4-0B-BC-0F
00-0A-E4-0B-D9-BE	00-0A-E4-0C-9D-0D
00-13-65-FE-F3-2C	00-13-65-FF-ED-2B
00-15-9B-FE-A4-66	00-15-9B-FF-24-B5
00-16-CA-00-00-00	00-16-CA-01-FF-FF
00-16-CA-F2-74-20	00-16-CA-F4-BE-0F
00-17-65-F6-94-C0	00-17-65-F7-38-CF
00-17-65-FD-00-00	00-17-65-FF-FF-FF
00-18-B0-33-90-00	00-18-B0-35-DF-FF
00-19-69-83-25-40	00-19-69-85-5F-FF

You can change these default MAC address ranges using the ACLI or EDM.

ADAC checks a MAC address against the supported ranges only after the MAC address is learned on the port. If you change the supported MAC address ranges, this has no effect on the previously learned MAC addresses. For example, if the address of a configured device is no longer in an ADAC range, the IP phone remains configured until its MAC address is aged out (by disconnecting the cable, for example) or until ADAC is disabled, either globally or on the port.

In a similar fashion, if the MAC address of an IP Phone—a MAC address that's not recognized by ADAC—is learned on a port and then is later added to the supported ranges, the IP Phone won't be detected and configured until the address is aged out or ADAC is disabled.

The maximum number of ranges that ADAC supports is 128.

Auto-Detection by LLDP (IEEE 802.1AB)

Auto-detection by LLDP extends the auto-detection that relies on MAC addresses. This feature allows devices identified as IP phones through LLDP to be detected by ADAC even if their MAC addresses are outside the list of ADAC MAC address ranges.

LLDP-based auto-detection supports a maximum of 16 devices per port.

ADAC and 802.1AB interoperability

With ADAC and 802.1AB interoperability, an IP phone configured with Avaya automatic QoS can update phone 802.1q priority and DSCP values based on Network Policy 802.1ab TLV values sent by the switch on an ADAC telephony port. The LLDP compliant IP phone then uses the received DSCP when sending voice traffic. Avaya Automatic QoS recognizes and prioritizes the traffic accordingly.

ADAC and 802.1AB interoperability is automatically enabled when Avaya automatic QoS, ADAC, and LLDP Network Policy TLV are enabled.

Note:

Because the ERS 2500 switches do not support user-configurable LLDP-MED network policies, LLDP implementation differs from the other Ethernet Routing Switch platforms. At LLDP default, the ERS 2500 switches tag voice traffic with a VID of ADAC Voice-VLAN ID instead of a VID of 0 (priority-tagged frames).

Auto-Configuration of Avaya IP Phones

The ADAC port participation can be set independently by enabling or disabling ADAC for particular ports.

When a new MAC address of an IP phone is learned on a port with ADAC enabled, ADAC immediately performs the auto-Configuration for that port (this operation is dependent on the configured ADAC operating mode and on whether other MAC addresses are learned on that port). This includes the required configuration of ports, VLANs, and QoS settings and involves minimal intervention by the user.

Auto-configuration is automatically removed or applied based on the port state, the state of the MAC addresses and the phones detected on the port.

The ports are polled every two seconds for their auto-configuration state and to see whether or not auto-configuration should be applied based on the current ADAC settings, both the global setting and the port setting. Auto-configuration will be applied on the port after the port is operational (operational state is enabled) and if one of these conditions is true:

- Op-mode = Untagged-Frames-Basic or Untagged-Frames-Advanced, at least one IP phone is detected on the port, and no non-IP phones are detected on the port
- Op-mode = Tagged-Frames and at least one IP phone is detected on the port

Auto-configuration is removed if any of these conditions becomes true:

- auto-detection becomes disabled on the port
- the ports operational state becomes disabled
- Op-mode = Untagged-Frames-Basic or -Advanced, and at least one non-IP device is detected on the port
- there are no IP phones detected on the port and the link is down.

If the link is still up but there are no IP phones on the port, auto-configuration is disabled after an aging period of about 90 seconds.

If all MAC addresses belonging to Avaya IP Phones on a port age out, the Auto-Configuration settings are removed from the port.

Chapter 8: VLAN configuration using ACLI

You configure and display VLANs using a variety of command modes. You can also enable or disable the automatic PVID feature.

Important:

For information about IGMP snooping and multicast groups, see *Avaya Ethernet Routing Switch 2500 Series Configuration—IP Multicast, NN47215-503*.

Displaying VLANs by type

Use the following procedure to display VLAN information.

Prerequisites

Log on to the Privileged EXEC mode in ACLI.

Procedure steps

Display VLANs by type:

```
show vlan [type {port | protocol}]
```

Note:

Enter `show vlan` to display all VLANs.

Table 7: Variable definitions

Variable	Value
type	Enter the type of VLAN. Values include: • port — show all port-based VLANs • protocol — show all protocol-based VLANs

Displaying VLAN settings per port

Use the following procedure to display VLAN settings associated with a port, including tagging information, PVID number, priority, and filtering information for tagged, untagged, and unregistered frames.

Procedure steps

To display VLAN settings associated with a port, use the following command from Privileged EXEC mode:

```
show vlan interface info [<portlist>]
```

Variable definitions

Variable	Value
<portlist>	Enter the list of ports for which you want the VLAN information, or enter *all* to display all ports.

Displaying port membership

Use the following procedure to display port membership in VLANs.

Procedure steps

To display port membership, use the following command from Privileged EXEC mode:

```
show vlan interface vids [<portlist>]
```

Variable definitions

Parameters and variables	Description
<portlist>	Enter the list of ports you want the VLAN information for, or enter all to display all ports.

Setting a management VLAN

Use the following procedure to set a management VLAN.

Procedure steps

To set a management VLAN, use the following command from Global Configuration mode:

```
vlan mgmt <1-4094>
```

Variable definitions

Variable	Value
<1-4094>	Enter the ID of the VLAN you want to serve as the management VLAN.

Deleting the management VLAN IP address

Use the following procedure to delete the management VLAN IP address.

Procedure steps

To delete the management VLAN IP address, use the following command from Global Configuration mode:

```
default ip address
```

Note:
This command will delete the management VLAN IP address from any mode.

Resetting the management VLAN

Use the following procedure to reset the management VLAN.

Procedure steps

To reset the management VLAN, use the following command from Global Configuration mode:

```
default vlan mgmt
```

Displaying VLAN ID

Use the following procedure to display a VLAN ID.

Procedure steps

To display a VLAN ID, use the following command from Privileged EXEC mode:

```
show vlan id <1-4094>
```

Variable definitions

Variable	Value
<1-4094>	Specifies the VLAN to be displayed.

Creating a VLAN

Use the following procedure to create a VLAN. You can create port-based or IPv6 protocol-based VLANs.

Procedure steps

To create a VLAN, use the following command from Global Configuration mode:

```
vlan create <1-4094> [name <line>] [type {port | protocol-ipv6Ether2}] {msti [<1-7> | cist]}
```

Important:
This command fails if the VLAN already exists.

Variable definitions

Variable	Value
<1-4094>	Enter the ID of the VLAN you want to create.
name <line>	Enter the new name you want for the VLAN.
type	Enter the type of VLAN. Values include: • port — port-based VLAN • protocol-ipv6Ether2 — IPv6 protocol-based VLAN
msti [<1-7> | cist]	This parameter is available only in MSTP mode. It associates the VLAN with either an MSTI instance or the CIST.

Deleting a VLAN

Use the following procedure to delete a VLAN.

Procedure steps

To delete a VLAN, use the following command from Global Configuration mode:

```
vlan delete <1-4094>
```

Variable definitions

Variable	Value
<1-4094>	Enter the ID of the VLAN to delete.

Deleting a VLAN - alternate method

Use the following procedure to delete a VLAN.

Procedure steps

To delete a VLAN, use the following command from Global Configuration mode:

```
no vlan <1-4094>
```

Variable definitions

Variable	Value
<1-4094>	Enter the ID of the VLAN you want to delete.

Configuring VLAN name

Use the following procedure to configure or change the name of a VLAN.

Procedure steps

To configure or change the name of a VLAN, use the following command from Global Configuration mode:

```
vlan name <1-4094> <line>
```

Variable definitions

Variable	Value
<1-4094>	Enter the ID of the VLAN for which you want to change the name.
<line>	Enter the new name you want for the VLAN.

Displaying VLAN Configuration Control settings

Use the following procedure to display current VLAN Configuration Control settings.

Procedure steps

To display VLAN Configuration Control settings, use the following command in Global Configuration mode:

```
show vlan configcontrol
```

Modifying VLAN Configuration Control settings

Use the following procedure to modify current VLAN Configuration Control settings. This command applies the selected option to all VLANs on the switch.

Procedure steps

1. To modify VLAN Configuration Control settings, use the following command from Global Configuration mode:

   ```
   vlan configcontrol <vcc_option>
   ```

2. To reset to the default value (strict), use the following command:

   ```
   default vlan configcontrol
   ```

Variable definitions

Variable	Value
<vcc_option>	This parameter denotes the VCC option to use on the switch. The valid values are: • automatic—Changes the VCC option to Automatic. • autopvid—Changes the VCC option to AutoPVID. • flexible—Changes the VCC option to Flexible. • strict—Changes the VCC option to Strict. This is the default VCC value.

Enabling automatic PVID

Use the following procedure to enable the automatic PVID feature. When auto PVID is active, a port that is assigned to a numbered VLAN has the same number for its PVID. For example, if the port belongs to VLAN 2, the port PVID is 2.

Procedure steps

To enable automatic PVID, use the following command from Global Configuration mode:

```
[no] auto-pvid
```

Use the **no** form of this command to disable.

Displaying automatic PVID status

Use the following procedure to display automatic PVID status.

Procedure steps

To display automatic PVID status, use the following command from User EXEC mode:

```
show auto-pvid
```

Configuring VLAN settings per port

Use the following procedure to configure VLAN settings per port.

Procedure steps

To configure VLAN settings per port, use the following command from Global Configuration mode:

```
vlan ports [<portlist>] [tagging{enable|disable|tagAll|
untagAll|tagPvidOnly|untagPvidOnly}][pvid <1-4094>] [filter-
untagged-frame {enable|disable}] [remarking {enable|disable}]
[priority <0-7>] [name <line>]
```

Variable definitions

Variable	Value
<portlist>	Enter the port numbers you want to configure for a VLAN.
tagging {enable| disable | tagAll | untagAll | tagPvidOnly | untagPvidOnly}	Specifies mode for PVID and non-PVID tagging.

Variable	Value
pvid <1-4094>	Associates the port with a specific VLAN.
filter-untagged-frame {enable|disable}	Enables or disables the port to filter received untagged packets.
filter-unregistered-frames{enable|disable}	Enables or disables the port to filter received unregistered packets.
remarking {enable|disable}	Enables or disables the priority of an incoming tagged frame using the priority for the ingress port.
priority <0-7>	Sets the port as a priority for the switch to consider as it forwards received packets.
name <line>	Enter the name you want for this port. **Important:** This option is available only if a single port is specified in the <portlist>.

Configuring VLAN members

Use the following procedure to add a port to or delete a port from a VLAN.

Procedure steps

To configure VLAN members, use the following command from Global Configuration mode:

```
vlan members [add|remove] <1-4094> <portlist>
```

Variable definitions

Variable	Value
add|remove	Adds a port to or removes a port from a VLAN. **Important:** If you omit this parameter, you set the exact port membership for the VLAN; the prior port membership of the VLAN is discarded and replaced by the new list of ports.

Variable	Value
<1-4094>	Specifies the target VLAN.
portlist	Enter the list of ports you wish to add, remove or assign to the VLAN.

MAC address table configuration using ACLI

This section describes how to view the contents of the MAC address forwarding database table, configure the age-out time for the addresses, and flush the MAC address table.

Important:

In certain situations, due to the hash algorithm used by switch to store MAC addresses into memory, some MAC addresses can not be learned.

Displaying the MAC address forwarding table using ACLI

Use the following procedure to display the current contents of the MAC address forwarding database table. You can now filter the MAC Address table by port number. The MAC address table can store up to 16000 addresses.

Prerequisites

Log on to the Privileged EXEC mode in ACLI.

Procedure steps

To display the MAC address forwarding table, use the following command:

```
show mac-address-table [vid <1-4094>] [aging-time] [address
<H.H.H | xx.xx.xx.xx.xx.xx | xx-xx-xx-xx-xx-xx>] [port
<portlist>]
```

Variable Definitions

Variable	Value
address <H.H.H | xx.xx.xx.xx.xx.xx | xx-xx-xx-xx-xx-xx>	Display a specific MAC address if it exists in the database. Enter the MAC address you want displayed using any of the 3 formats.
aging-time	Display the time in seconds after which an unused entry is removed from the forwarding database.
port <portlist>	Specify ports.
vid <1-4094>	Enter the ID of the VLAN for which you want to display the forwarding database. Default is to display the management VLAN's database.

Configuring aging time for unseen MAC addresses using ACLI

Use the following procedure to configure the time during which the switch retains unseen MAC addresses.

Prerequisites

Log on to the Global Configuration mode in ACLI.

Procedure steps

To configure aging time, use the following command:

```
mac-address-table aging-time <10-1 000 000>
```

Variable Definitions

Variable	Value
<10-1 000 000>	Enter the aging time in seconds that you want for MAC addresses before they expire.

Configuring aging time for unseen MAC addresses to default using ACLI

Use the following procedure to set the aging time for MAC addresses to 300 seconds.

Prerequisites

Log on to the Global Configuration mode in ACLI.

Procedure steps

To set again time to default (300 seconds), use the following command:

```
default mac-address-table aging-time
```

Flushing the MAC address table using ACLI

Flush the MAC address table to clear all addresses in the MAC address table.

Prerequisites

Log on to the Privileged EXEC mode in ACLI.

Procedure steps

To flush the MAC address table, use the following command:

```
clear mac-address-table
```

Flushing a VLAN MAC address table using ACLI

Flush the MAC address table for a VLAN to clear the MAC addresses for a specific VLAN.

Prerequisites

Log on to the Privileged EXEC mode in ACLI.

Procedure steps

To flush the MAC address table for a specific VLAN, use the following command:

```
clear mac-address-table interface vlan <1-4094>
```

Variable definition

Use the information in the following table to flush the MAC address table for a VLAN.

Variable	Value
1-4094	Specify the VLAN for which you want to be flush the MAC addresses.

Flushing a FastEthernet interface MAC address table using ACLI

Flush the MAC address table for a FastEthernet interface to clear the MAC addresses for specified ports. This command does not flush the addresses learned on the trunk.

Prerequisites

Log on to the Privileged EXEC mode in ACLI.

Procedure steps

To flush the MAC address table on a FastEthernet interface, use the following command.

```
clear mac-address-table interface FastEthernet <LINE>
```

Variable definition

Use the information in the table to flush the MAC address table for a FastEthernet interface.

Variable	Value
LINE	Specifies the list of ports, in the slot/port format, for which you want to flush the MAC addresses .

Flushing the MAC address table for a trunk using ACLI

Flush the MAC address table for a trunk to clear the MAC addresses for the specified trunk. This command flushes only addresses that are learned on the trunk.

Prerequisites

Log on to the Privileged EXEC mode in ACLI.

Procedure steps

To clear the MAC address table on a trunk, use the following command:

```
clear mac-address-table interface mlt <1-32>
```

Variable definition

Use the information in the table to complete this procedure.

Variable	Value
1-32	Specifies the Trunk for which you want to flushed the MAC addresses.

Flushing a single address from the MAC address table using ACLI

Flush a single address from the MAC address table to clear one MAC address from the MAC address table.

Prerequisites

Log on to the Privileged EXEC mode in ACLI.

Procedure steps

To flush a single MAC address, use the following command:

```
clear mac-address-table address <H.H.H | xx.xx.xx.xx.xx.xx>
```

Variable definition

Use the information in the table to complete this procedure.

Variable	Value
H.H.H | xx.xx.xx.xx.xx.xx | xx-xx-xx-xx-xx-xx	Specifies the MAC address to clear.

Chapter 9: STP configuration using ACLI

This chapter describes how to configure the Spanning Tree Protocol using the command line interface (ACLI).

Using spanning tree

You can use the ACLI to configure a spanning tree, to add or remove VLANs to the spanning tree, and to configure the usual spanning tree parameters and FastLearn.

For detailed information about spanning tree parameters, Spanning Tree Groups, and configuration guidelines, see Spanning Tree Protocol Fundamentals on page 27.

Displaying spanning tree configuration information

Use the following procedure to display spanning tree configuration information that is specific to either the spanning tree group or to the port.

Procedure steps

To display spanning tree configuration information, use the following command from Privileged EXEC mode:

```
show spanning-tree {config|port|op-mode|cost-calc-mode}
```

Variable definitions

Variable	Value
config	Displays spanning tree configuration.
port	Displays spanning tree status of each port.
op-mode	Displays the spanning tree operational mode.
cost-calc-mode	Displays pathcost type.

Setting path cost calculation

Use the following procedure to set the path cost calculation mode for the Spanning Tree Group.

Procedure steps

To set the path cost calculation, use the following command from Global Configuration mode:

```
default spanning-tree cost-calc-mode
```

Configuring STG parameters

Use the following procedure to configure Spanning Tree Group parameters.

Procedure steps

1. To configure STG parameters, use the following command from Global Configuration mode:

   ```
   spanning-tree [forward-time <4-30>] [hello-time <1-10>][max-age <6-40>][priority {0*0000 | 0*1000| 0*2000 | 0*3000 | ... | 0*E000 | 0*F000}]
   ```

2. To reset to default, use the following command:

   ```
   default spanning-tree [forward-time] [hello-time][max-age] [priority]
   ```

Variable definitions

Variable	Value
forward-time <4-30>	Enter the forward time of the STG in seconds; the range is 4 -- 30, and the default value is 15.
hello-time <1-10>	Enter the hello time of the STG in seconds; the range is 1 --10, and the default value is 2.
[max-age <6-40>]	Enter the max-age of the STG in seconds; the range is 6 -- 40, and the default value is 20.

Variable	Value
[priority {0*0000 | 0*1000| 0*2000 | 0*3000 | ... | 0*E000 | 0*F000}]	Sets the spanning tree priority (in Hex); if 802.1T compliant, this value must be a multiple of 0x1000.

Configuring STG operation mode

Use the following procedure to set the operation mode for the Spanning Tree Group.

Procedure steps

To set the operation mode, use the following command from Global Configuration mode:

```
spanning-tree op-mode {mstp|rstp|stpg}
```

Warning:

To prevent the stack from losing its configuration, multiple power cycling (hard resets) is not recommended after alternately changing spanning-tree operation mode.

Configuring STP for ports

Use the following procedure to configure Spanning Tree Protocol for ports.

Procedure steps

To configure STP, use the following command from Interface Configuration mode:

```
spanning-tree [port <portlist>] [learning {disable|normal|
fast}] [cost <1-65535>] [priority <0-255>]
```

Variable definitions

Variable	Value
port <portlist>	Enables spanning tree for the specified port or ports; enter the port or ports you want enabled for spanning tree.

Variable	Value
	Important: If you omit this parameter, the system uses the port number you specified after you issued the interface command.
learning {disable|normal|fast}	Specifies the STP learning mode: • disable—disable spanning tree on the port • normal—normal learning mode • fast—FastLearn mode
cost <1-65535>	Enter the path cost of the spanning tree; range is 1 to 65535.
priority <0-255>	Enter the priority value of the spanning tree; range is 0 to 255.

Configuring STP values per port

Use the following procedure to set the spanning tree values for the ports within the specified spanning tree group to the factory default settings.

Procedure steps

To set STP values, use the following command from Interface Configuration mode:

```
default spanning-tree [port <portlist>] [learning] [cost] [priority]
```

Variable definitions

Variables	Value
port <portlist>	Enables spanning tree for the specified port or ports; enter the port or ports you want set to factory spanning tree default values. **Important:** If you omit this parameter, the system uses the port number you specified after you issued the **`interface`** command.
learning	Sets the spanning tree learning mode to factory default value. Default value for learning is normal mode.
cost	Sets the path cost to factory default value. Default value for path cost depends on the type of port.
priority	Sets the priority to factory default value.

Variables	Value
	Default value for the priority is 0x8000.

Configuring STP port mode using ACLI

Configure Spanning Tree port mode to enable a port to maintain STP membership when the port is moved from one VLAN to another.

Prerequisites

Log on to the Global Configuration mode in ACLI.

Procedure steps

To configure Spanning Tree port mode, use the following command:

```
spanning-tree port-mode {auto | normal}
```

Variable definitions

Variable	Value
`auto`	Specifies automatic STP port mode.
`normal`	Specifies normal STP port mode.

Enabling STP 802.1d compliance mode using ACLI

Enable STP 802.1d compliance mode to ensure that STP conforms to the IEEE 802.1d standard.

Prerequisites

Log on to the Global Configuration mode in ACLI.

Procedure steps

Enable STP 802.1d compliance mode by using the following command:

```
spanning-tree 802dot1d-port-compliance enable
```

Disabling STP 802.1d compliance mode using ACLI

Disable STP 802.1d compliance mode to release STP from conforming to the IEEE 802.1d standard.

Prerequisites

Log on to the Global Configuration mode in ACLI.

Procedure steps

Disable STP 802.1d compliance mode by using the following command:

```
no spanning-tree 802dot1d-port-compliance enable
```

Disabling STP for ports

Use the following procedure to disable STP for ports in a specific STG.

Procedure steps

To disable STP, use the following command from Interface Configuration mode:

```
no spanning-tree [port <portlist>]
```

Variable definitions

Variable	Value
port <portlist>	Disables spanning tree for the specified port or ports; enter port or ports you want enabled for STP. **Important:** If you omit this parameter, the system uses the port number you specified after you issued the **`interface`** command.

Using Advanced Spanning Tree

The Advanced Spanning Tree Protocol (ASTP) application comprises Rapid Spanning Tree Protocol (RSTP) and Multi Spanning Tree Protocol (MSTP). You can configure the RSTP and the MSTP applications.

Displaying RSTP configuration details

Use the following procedure to display the RSTP related bridge-level configuration details.

Procedure steps

To display RSTP information, use the following command from Privileged EXEC mode:

```
show spanning-tree rstp config
```

Displaying RSTP bridge statistics

Use the following procedure to display the RSTP related bridge-level statistics.

Procedure steps

To display RSTP statistics, use the following command from Privileged EXEC mode:

```
show spanning-tree rstp statistics
```

Displaying RSTP status information

Use the following procedure to display the RSTP related status information for the selected bridge.

Procedure steps

To display RSTP status information, use the following command from Privileged EXEC mode:

```
show spanning-tree rstp status
```

Displaying RSTP port configuration details

Use the following procedure to display RSTP related port-level configuration details.

Procedure steps

To display RSTP configuration details, use the following command from Privileged EXEC mode:

```
show spanning-tree rstp port config[<portlist>]
```

Display RSTP port statistics

Use the following procedure to display RSTP related port-level statistics.

Procedure steps

To display RSTP statistics, use the following command from Privileged EXEC mode:

```
show spanning-tree rstp port statistics <portlist>
```

Displaying RSTP status per port

Use the following procedure to display the RSTP related status information for the selected port.

Procedure steps

To display RSTP status, use the following command from Privileged EXEC mode:

```
show spanning-tree rstp port status [<portlist>]
```

Configuring RSTP parameters

Use the following procedure to set the RSTP parameters, which include forward delay, hello time, maximum age time, default pathcost version, bridge priority, transmit hold count, and version for the bridge.

Procedure steps

To set RSTP parameters, use the following command from Global Configuration mode:

```
spanning-tree rstp [port <portlist>] [cost <1 - 200000000>]
[edge-port {false | true}][learning {disable | enable}][p2p
{auto | force-false | force-true}][priority {00 | 10 | … | F0}]
[protocol-migration {false | true}]
```

Variable definitions

Parameters and variables	Description
port <portlist>	Filter on list of ports.
cost <1 - 200000000>	Sets the RSTP pathcost on the single or multiple ports; default is 200000.
edge-port {false | true}	Indicates whether the single or multiple ports should be assumed to be edge port. This parameter sets the Admin value of edge port status; default is false.
learning {disable | enable}	Enables or disables RSTP on the single or multiple ports; default is enable.

Parameters and variables	Description
p2p {auto | force-false | force-true}	Indicates whether the single or multiple port should be treated as a point-to-point link or not. This command sets the Admin value of P2P Status; default is force-true.
priority {00 | 10 |… | F0}	Sets the RSTP port priority on the single or multiple port; default is 80.
protocol-migration {false | true}	Forces the single or multiple port to transmit RSTP BPDUs when set true, while operating in RSTP mode; default is false.

Configuring RSTP parameters

Use the following procedure to set the RSTP parameters, which include pathcost, edge-port indicator, learning mode, point-to-point indicator, priority, and protocol migration indicator on the single or multiple ports.

Procedure steps

To configure parameters, use the following command from Interface Configuration mode:

```
spanning-tree rstp port <portlist>
```

Variable definitions

Parameters and variables	Description
port <portlist>	Filter on list of ports.

Displaying MSTP related information

Use the following procedure to display the MSTP related bridge-level, VLAN and region information.

Procedure steps

To display MSTP related information, use the following command from Privileged EXEC mode:

```
show spanning-tree mstp config
```

Displaying MSTP related statistics

Use the following procedure to display MSTP related bridge-level statistics.

Procedure steps

To display MSTP statistics, use the following command from Privileged EXEC mode:

```
show spanning-tree mstp statistics
```

Displaying MSTP status information

Use the following procedure to display the MSTP related status information known by the selected bridge.

Procedure steps

To display MSTP status information, use the following command from Privileged EXEC mode:

```
show spanning-tree mstp status
```

Displaying MSTP Cist port information

Use the following procedure to display the Multi Spanning Tree Protocol (MSTP) Cist Port information maintained by every port of the Common Spanning Tree.

Procedure steps

To display information, use the following command from Privileged EXEC mode:

```
show spanning-tree mstp port config [<portlist>]
```

Important:

In MSTP, if the Regional Root changes, the change does not display correctly when entering the **`show spanning-tree mstp port config`** command. In the

command output, the Cist Port Regional Root field does not display the correct Regional Root.

Variable definitions

Parameters and variables	Description
<portlist>	Enter a list or range of port numbers.

Displaying MSTP Cist port statistics

Use the following procedure to display the Multi Spanning Tree Protocol (MSTP) Cist Port statistics maintained by every port.

Procedure steps

To display statistics, use the following command from Privileged EXEC mode:

```
show spanning-tree mstp port statistics [<portlist>]
```

Variable definitions

Parameters and variables	Description
<portlist>	Enter a list or range of port numbers.

Displaying MSTP bridge and VLAN information

Use the following procedure to display the Multi Spanning Tree Protocol (MSTP) instance-specific bridge and VLAN information.

Procedure steps

To display information, use the following command from Privileged EXEC mode:

```
show spanning-tree mstp msti config <1 - 7>
```

Variable definitions

Parameters and variables	Description
<1 - 7>	Filter on MSTP instance.

Displaying MSTP bridge statistics

Use the following procedure to display the Multi Spanning Tree Protocol (MSTP) instance-specific bridge statistics.

Procedure steps

To display statistics, use the following command from Privileged EXEC mode:

```
show spanning-tree mstp msti statistics <1 - 7>
```

Variable definitions

Parameters and variables	Description
<1 - 7>	Filter on MSTP instance.

Displaying MSTP port information

Use the following procedure to display the Multi Spanning Tree Protocol (MSTP) instance-specific to port information.

Procedure steps

To display port information, use the following command from Privileged EXEC mode:

```
show spanning-tree mstp msti port config <1 - 7> [<portlist>]
```

Variable definitions

Parameters and variables	Description
<1 - 7>	Filter on MSTP instance.
[<portlist>]	Enter a list or range of port numbers.

Displaying MSTP port statistics

Use the following procedure to display the Multi Spanning Tree Protocol (MSTP) instance-specific to port statistics.

Procedure steps

To display statistics, use the following command from Privileged EXEC mode:

```
show spanning-tree mstp msti port statistics <1 - 7>
[<portlist>]
```

Variable definitions

Parameters and variables	Description
<1 - 7>	Filter on MSTP instance.
[<portlist>]	Enter a list or range of port numbers.

Configuring MSTP parameters for Cist bridge

Use the following procedure to configure the MSTP parameters which includes maximum hop count, maximum number of instances allowed, forward delay time, hello time, maximum age time, default pathcost version, priority, transmit hold count, and version for the Cist Bridge.

Procedure steps

To configure parameters, use the following command from Global Configuration mode:

```
spanning-tree mstp [max-hop <600 - 4000>] [forward-time <4 -
30>] [max-age <6 - 40>] [pathcost-type {bits16 | bits32}]
[priority {0000 | 10000 | 20000 | … | F0000}] [tx-hold count <1
- 10>] [version {stp-compatible | rstp| mstp}] [add-vlan
<1-4094>] [remove-vlan <1-4094>] [msti <1-7>] [region {config-
id-sel|region-name|region-version}]
```

Variable definitions

Variable	Value
max-hop <600 - 4000>	Sets the MSTP maximum hop count; default is 2000
forward-time <4 - 30>	Sets the MSTP forward delay for the Cist Bridge in seconds; default is 15
max-age <6 - 40>	Sets the MSTP maximum age time for the Cist Bridge in seconds; default is 20
pathcost-type {bits16 | bits32}	Sets the MSTP default pathcost version; default is bits32
priority {0000 | 10000| 20000 … | F000}	Sets the MSTP bridge priority for the Cist Bridge; default is 8000
tx-holdcount<1 - 10>	Sets the MSTP Transmit Hold Count; default is 3
version {stp-compatible | rstp | mstp}	Sets the MSTP version for the Cist Bridge; default is mstp
add-vlan	Adds a VLAN to the CIST bridge.
remove-vlan	Removes a VLAN from the CIST bridge.
msti	Changes MSTP instance-specific configuration.
region	Changes MSTP region configuration.

Configuring MSTP parameters for Common Spanning Tree

Use the following procedure to configure the MSTP parameters which includes pathcost, hello time, edge-port indicator, learning mode, point-to-point indicator, priority, and protocol migration indicator on the single or multiple port for the Common Spanning Tree.

Procedure steps

To configure parameters, use the following command from Interface Configuration mode:

```
spanning-tree mstp [port <portlist>] [cost <1 - 200000000>]
[edge-port {false | true}][hello-time <1 - 10>] [learning
```

```
{disable | enable}][p2p {auto | force-false | force-true}]
[priority {00 | 10 | … | F0}] [protocol-migration {false |
true}]
```

Variable definitions

Variable	Value
port <portlist>	Enter a list or range of port numbers.
cost <1 - 200000000>	Sets the MSTP pathcost on the single or multiple port; default is 200000
hello-time <1 - 10>	Sets the MSTP hello time on the single or multiple port for the Common Spanning Tree; default is 2
edge-port {false | true}	Indicates whether the single or multiple port should be assumed to be edge port or not. This parameter sets the Admin value of edge port status; default is false
learning {disable | enable}	Enables or disables MSTP on the single or multiple port; default is enable
p2p {auto | force-false | force-true}	Indicates whether the single or multiple port should be treated as a point-to-point link or not. This command sets the Admin value of P2P Status; default is force-true
priority {00 | 10 |… | F0}	Sets the MSTP port priority on the single or multiple port; default is 80
protocol-migration {false | true}	Forces the single or multiple port to transmit MSTP BPDUs when set true, while operating in MSTP mode; default is false

Configuring MSTP region parameters

Use the following procedure to configure the MSTP parameters which includes config ID selector, region name, and region version.

Procedure steps

To configure parameters, use the following command from Global Configuration mode:

```
spanning-tree mstp region [config-id-sell <0 - 255>] [region-
name <1 - 32 chars>][region-version <0 - 65535>]
```

Variable definitions

Variable	Value
[config-id-sel <0 - 255>]	Sets the MSTP config ID selector; default is 0
[region-name <1 - 32 chars>]	Sets the MSTP region name, default is the MAC address of the switch.
[region-version <0 - 65535>]	Sets the MSTP region version; default is 0

Configuring MSTP MSTI bridge parameters

Use the following procedure to configure the MSTP parameters which includes forward delay time, hello-time, max hop count, priority, and VLAN mapping for the bridge instance.

Procedure steps

To configure parameters, use the following command from Global Configuration mode:

```
spanning-tree mstp msti<1 - 7>[priority{0000|1000|…|F000}]
[add-vlan <vid>][remove-vlan <vid>][enable]
```

Variable definitions

Variable	Value
<1 - 7>	Filter on MSTP instance.
priority {0000 | 1000 |… | F000}	Sets the MSTP priority for the bridge instance; default is 8000
add-vlan <1 - 4094>	Maps the specified vlan and MSTP bridge instance
remove-vlan <1 - 4094>	Unmaps the specified vlan and MSTP bridge instance
enable	Enables the MSTP bridge instances

Configuring MSTP MSTI port parameters

Use the following procedure to configure the MSTP parameters which includes MSTP port pathcost, learning mode, and priority on the single or multiple port for the bridge instance.

Procedure steps

To configure parameters, use the following command from Interface Configuration mode:

```
spanning-tree mstp msti <1 - 7> [port <portlist>] [cost <1 -
200000000>][learning {disable | enable}][priority {00 | 10 | …
| F0}]
```

Variable definitions

Variable	Value
<1 - 7>	Filter on MSTP instance.
port <portlist>	Enter a list or range of port numbers.
cost <1 - 200000000>	Sets the MSTP port pathcost on the single or multiple port for the bridge instance; default is 200000
learning {disable | enable}	Enables or disables MSTP on the single or multiple port for the bridge instance; default is enable
priority {00 | 10 |… | F0}	Sets the MSTP port priority on the single or multiple port for the bridge instance; default is 80

Deleting a MSTP bridge

Use the following procedure to delete a MSTP bridge-instance.

Procedure steps

To delete a bridge instance, use the following command from Interface Configuration mode:

```
no spanning-tree mstp msti <1 - 7>
```

Enabling a MSTP bridge

Use the following procedure to enable a MSTP bridge instance.

Procedure steps

To enable a bridge instance, use the following command from Interface Configuration mode:

```
[no] spanning-tree mstp msti <1 - 7> enable
```

Use the no form of this command to disable.

Configuring STP BPDU filtering

Use the following procedure to configure STP BPDU Filtering on a port. This command is available in all STP modes (STPG, RSTP, and MSTP).

Procedure steps

1. To configure STP BPDU filtering, use the following command from Interface Configuration mode:

   ```
   spanning-tree bpdu-filtering [port <portlist>] [enable]
   [timeout <10-65535 | 0> ]
   ```

2. To return to default values, use the following command:

   ```
   default spanning-tree bpdu-filtering [port <portlist>]
   [enable] [timeout]
   ```

3. To disable, use the following command:

   ```
   no spanning-tree bpdu-filtering [port <portlist>] [enable]
   ```

4. To display the status of parameters, use the following command:

   ```
   show spanning-tree bpdu-filtering [<interface-type>][port
   <portlist>]
   ```

Variable definitions

Parameter	Description
port <portlist>	Specifies the ports affected by the command.
enable	Enables STP BPDU Filtering on the specified ports. The default value is disabled.
timeout <10-65535 | 0 >	When BPDU filtering is enabled, this indicates the time (in seconds) during which the port remains disabled after it receives a BPDU. The port timer is disabled if this value is set to 0. The default value is 120 seconds.

Chapter 10: Multi-Link Trunking configuration using ACLI

This chapter describes how to configure Multi-Link Trunking (MLT), Link Aggregation Groups (LAG), Distributed Multi-Link Trunking (DMLT) and Distributed Line Aggregation LAG (802.3ad LACP) using the ACLI.

Configuring Multi-Link Trunking

Displaying MLT configuration

Use the following procedure to display Multi-Link Trunking (MLT) configuration and utilization.

Procedure steps

To display MLT configuration and utilization, use the following command from Global Configuration mode.

```
show mlt [utilization <1-6>]
```

Variable definitions

Variable	Value
utilization <1-6>	Displays the utilization of the specified enabled MLT(s) in percentages.

Configuring a Multi-Link Trunk

Use the following procedure to configure a Multi-Link Trunk.

Procedure steps

To configure a Multi-Link Trunk, use the following command from Global Configuration mode.

```
mlt <id> [name <trunkname>] [enable|disable] [member
<portlist>][learning {disable|fast|normal}] [loadbalance
<advance|basic>][bpdu{all-ports|single-port}]
```

Important:
An MLT must be disabled when you are adding ports.

Variable definitions

Variable	Value
id	Enter the trunk ID; range is 1 to 6.
name <trunkname>	Specifies a text name for the trunk; enter up to 16 alphanumeric characters.
enable|disable	Enables or disables the trunk.
member <portlist>	Enter the ports that you want as members of the trunk.
learning <disable|fast|normal>	Sets STP learning mode.
loadbalance <advance|basic>	Specifies MLT load balancing mode. Advance mode uses IP based load balancing. Basic mode uses MAC based load balancing.
bpdu{all-ports|single-port}	Set BPDU send/receive mode.

Disabling a Multi-Link Trunk

Use the following procedure to delete a specific Multi-Link Trunk (MLT) or all configured MLTs.

Procedure steps

1. To delete a specific MLT, use the following command from Global Configuration mode.

```
no mlt [<id>]
```

2. To delete all configured MLTs, use the following command from Global Configuration mode.

```
no mlt
```

Variable definitions

Variable	Value
<id>	Specifies the ID of the MLT you want to delete.

Configuring MLT whole trunk using ACLI

Use this procedure to configure the shutdown of all ports in the MLT. This procedure enables or disables the MLT whole trunk feature.

Prerequisites

Log on to the Global configuration mode in ACLI.

Procedure steps

Configure the shutdown of all ports in the MLT:

```
[no] mlt shutdown-ports-on-disable enable
```

Variable definitions

Variable	Value
[no]	Disables the MLT whole trunk feature.

Displaying the MLT whole trunk status using ACLI

Use this procedure to display the current MLT whole trunk mode.

Prerequisites

Log on to the Priviledged EXEC mode in ACLI.

Procedure steps

Display the current MLT whole trunk mode:

```
show mlt shutdown-ports-on-disable
```

Job aid

The following shows example outputs for the **show mlt shutdown-ports-on-disable** command.

`show mlt shutdown-ports-on-disable`

`Trunk loop prevention is disabled`— MLT whole trunk feature is disabled (default).

`show mlt shutdown-ports-on-disable`

`Trunk loop prevention is enabled`— MLT whole trunk feature is enabled.

Configuring Link Aggregation Group

Configuring LACP system priority

Use the following procedure to set a system priority for LACP.

Procedure steps

To set a system priority for LACP, use the following command from Global Configuration mode.

```
lacp system-priority [0-65535]
```

Variable definitions

Variable	Value
[0-65535]	Enter a system priority for LACP; range is 0 to **65535**.

Resetting LACP system priority to default

Use the following procedure to set the system priority for LACP as the default value of 32768.

Procedure steps

To set the system priority for LACP to default, use the following command from Global Configuration mode.

```
default lacp system-priority
```

Configuring LACP port mode

Use the following procedure to set the mode for an LACP port.

Procedure steps

To set the mode for an LACP port, use the following command from Interface Configuration mode.

```
lacp mode [port <portlist>] {off|passive|active}
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want to set LACP the mode.
port {off | passive | active}	Sets the LACP mode for the specified port to off, passive, or active; if port mode is selected as Passive or Active, port is ready to participate in LACP

Resetting LACP port mode to default

Use the following procedure to place an LACP port in the default off mode.

Procedure steps

To place an LACP port in default mode, use the following command from Interface Configuration mode.

```
default lacp mode [port <portlist>]
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports that you want to set in the LACP off mode by default.

Enabling LACP aggregation

Use the following procedure to enable LACP aggregation on the specified port(s).

Procedure steps

To enable LACP aggregation on specified ports, use the following command from Interface Configuration mode.

```
lacp aggregation [port <portlist>] enable
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want to enable LACP aggregation.

Removing LACP aggregation for ports

Use the following procedure to remove LACP aggregation on the specified port(s).

Procedure steps

To remove LACP aggregation for ports, use the following command from Interface Configuration mode.

```
no lacp aggregation [port <portlist>] enable
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want to disable LACP aggregation.

Disabling LACP for ports

Use the following procedure to disable LACP aggregation on the specified port(s) by default.

Procedure steps

To disable LACP aggregation for ports, use the following command from Interface Configuration mode.

```
default lacp aggregation [port <portlist>] enable
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want to disable LACP aggregation by default.

Assigning a key value to a port

Use the following procedure to assign a key value for the specified port(s).

Procedure steps

1. To assign a key value, use the following command from Interface Configuration mode.

   ```
   default lacp key [port <portlist>]
   ```

2. To set the LACP key to the default value (1), enter the following command:

```
default lacp key [port <portlist>]
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want to assign an LACP key value.
<1-4095>	Enter an LACP key value for the port; range is 1 to 4095.

Assigning LACP priority for ports

Use the following procedure to set an LACP priority for the specified port(s).

Procedure steps

To set LACP priority for ports, use the following command from Interface Configuration mode.

```
lacp priority [port <portlist>] <0-65535>
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want to set LACP priority.
<0-65535>	Enter a priority number for the port; range is 0 to 65535.

Setting LACP priority to default

Use the following procedure to set the LACP priority for the specified port(s) as the default value of 128.

Procedure steps

To set LACP priority to default, use the following command from Interface Configuration mode.

```
default lacp priority [port <portlist>]
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want to set the default LACP priority of 128.

Configuring LACP timeout

Use the following procedure to set an LACP timeout for the specified port(s).

Procedure steps

To configure LACP timeout, use the following command from Interface Configuration mode.

```
lacp timeout-time [port <portlist>] {short | long}
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want to set an LACP timeout.
port {short | long}	Sets the a short or long LACP timeout for the port. The long timeout is 90 seconds and short timeout is 3 seconds.

Configuring long LACP timeout for ports

Use the following procedure to set a long LACP timeout for the specified port(s) by default.

Procedure steps

To set a long LACP timeout for ports, use the following command from Interface Configuration mode.

```
default lacp timeout-time [port <portlist>]
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want to set a long LACP timeout by default.

Displaying LACP information

Use the following procedure to display LACP information for the entire system.

Procedure steps

To display LACP information, use the following command from Interface Configuration mode.

```
show lacp system
```

Displaying LACP aggregator information

Use the following procedure to display LACP aggregator information.

Procedure steps

To display LACP aggregator information, use the following command from Global Configuration mode.

```
show lacp aggr [<1-65535>]
```

Variable definitions

Variable	Value
<1-65535>	Enter the aggregator ID.

Displaying LACP port information

Use the following procedure to display LACP port information.

Procedure steps

To display LACP port information, use the following command from Interface Configuration mode.

```
show lacp port [<portlist>]
```

Important:
The output of the **`show vlacp port`** command will display "A" or "I" for port type. A=Aggregatable and I=Individual.

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want information.

Displaying LACP port debug information

Use the following procedure to display LACP port debug information.

Procedure steps

To display LACP port debug information, use the following command from Interface Configuration mode:

```
show lacp debug member [portlist]
```

The command can display the following terms.

LACP Receiving State:

- Current: Rx information is valid
- Expired: Rx information is invalid.
- Defaulted: Rx machine is defaulted.
- Initialized: Rx machine is initializing.
- LacpDisabled: LACP is disabled on this port.
- PortDisabled: Port is disabled.

Selection State:

- Detached: port is not attached to any aggregator.
- Waiting: port is waiting to attach to an aggregator.
- Attached: port is attached to an Aggregator.
- Ready: port is ready to Tx and Rx.

Variable definitions

Variable	Value
port [portlist]	Enter the ports for which you want debug information.

Displaying LACP port statistics

Use the following procedure to display LACP port statistics.

Procedure steps

To display LACP port statistics, use the following command from Interface Configuration mode:

```
show lacp stats [port <portlist>]
```

Variable definitions

Variable	Value
port <portlist>	Enter the ports for which you want statistics.

Clearing LACP port statistics

Use the following procedure to clear LACP port statistics.

Procedure steps

To clear LACP port statistics, use the following command from Interface Configuration mode:

```
lacp clear-stats [port <portlist>]
```

Configuring VLACP using the ACLI

You can use the ACLI to configure VLACP parameters.

Note:
When you set VLACP parameters for a trunk port, the settings are applied to all trunk members.

Enabling VLACP

Use the following procedure to globally enable VLACP for the device.

Procedure steps

To enable VLACP, use the following command from Global Configuration mode:

```
vlacp enable
```

Configuring multicast MAC address for VLACP

Use the following procedure to set the multicast MAC address used by the device for VLACPDUs.

Procedure steps

To set the multicast MAC address, use the following command from Global Configuration mode:

```
vlacp macaddress <macaddress>
```

Configuring VLACP on a port

Use the following procedure to configure VLACP parameters on a port.

Procedure steps

To configure VLACP parameters on a port, use the following command from Interface Configuration mode:

```
vlacp port <slot/port> [enable | disable] [timeout <long/
short>] [fast-periodic-time <integer>] [slow-periodic-time
<integer>] [timeout-scale <integer>] [funcmac-addr <mac>]
[ethertype <hex>]
```

Variable definitions

Variable	Value
<slot/port>	Specifies the slot and port number.
enable|disable	Enables or disables VLACP.
timeout <long/short>	Specifies whether the timeout control value for the port is a long or short timeout.

Variable	Value
	• long sets the port timeout value to: (timeout-scale value) × (slow-periodic-time value). • short sets the port's timeout value to: (timeout-scale value) × (fast-periodic-time value). For example, if the timeout is set to short while the timeout-scale value is 3 and the fast-periodic-time value is 400 ms, the timer expires after 1200 ms. Default is long.
fast-periodic-time <integer>	Specifies the number of milliseconds between periodic VLACPDU transmissions using short timeouts. The range is 400-20000 milliseconds. Default is 500.
slow-periodic-time <integer>	Specifies the number of milliseconds between periodic VLACPDU transmissions using long timeouts. The range is 10000-30000 milliseconds. Default is 30000.
timeout-scale <integer>	Sets a timeout scale for the port, where timeout = (periodic time) × (timeout scale). The range is 1-10. Default is 3. Note: With VLACP, a short interval exists between a port transmitting a VLACPDU and the partner port receiving the same VLACPDU. However, if the timeout-scale is set to 1, the port timeout value does not take into account the normal travel time of the VLACPDU. The port expects to receive a VLACPDU at the same moment the partner port sends it. Therefore, the delayed VLACPDU results in the link being blocked, and then enabled again after the packet arrives. To prevent this scenario from happening, set the timeout-scale to a value larger than 1.
funcmac-addr <mac>	Specifies the address of the far-end switch/stack configured to be the partner of this switch/stack. If none is configured, any VLACP-enabled switch communicating with the local switch through VLACP PDUs is considered to be the partner switch. Note: VLACP has only one multicast MAC address, configured using the vlacp macaddress command, which is the Layer 2 destination address used for the VLACPDUs. The port-specific funcmac-addr parameter does not specify a multicast MAC address, but instead

Variable	Value
	specifies the MAC address of the switch or stack to which this port is sending VLACPDUs. You are not always required to configure funcmac-addr. If not configured, the first VLACP-enabled switch that receives the PDUs from a unit assumes that it is the intended recipient and processes the PDUs accordingly. If you want an intermediate switch to drop VLACP packets, configure the funcmac-addr parameter to the desired destination MAC address. With funcmac-addr configured, the intermediate switches do not misinterpret the VLACP packets.
ethertype <hex>	Sets the VLACP protocol identification for this port. Defines the ethertype value of the VLACP frame. The range is 8101-81FF. Default is 8103.

Disabling VLACP globally

Use the following procedure to globally disable VLACP for the device.

Procedure steps

To globally disable VLACP for a device, use the following command from Global Configuration mode:

```
no vlacp enable
```

Resetting VLACP MAC address value

Use the following procedure to reset the multicast MAC address used by the device for VLACPDUs to the default value (01:80:c2:00:11:00).

Procedure steps

To reset the MAC address, use the following command from Global Configuration mode:

```
no vlacp macaddress
```

Disabling VLACP on a port

Use the following procedure to disable VLACP on the port.

Procedure steps

To disable VLACP on a port, use the following command from Global Configuration mode:

```
no vlacp <slot/port> [enable] [funcmac-addr]
```

Variable definitions

Variable	Value
<slot/port>	Specifies the slot and port number.
enable	Disables VLACP on the specified port.
funcmac-addr	Sets the funcmac-addr parameter to the default value.

Displaying VLACP status

Use the following procedure to display the status of VLACP on the switch.

Procedure steps

To display the status of VLACP, use the following command from Privileged EXEC mode:

```
show vlacp
```

Displaying VLACP configuration for a port

Use the following procedure to display VLACP configuration details for a port or list of ports.

Procedure steps

To display VLACP configuration details per port, use the following command from Privileged EXEC mode:

```
show vlacp interface <slot/port>
```

where `<slot/port>` specifies a port or list of ports.

Among other properties, the `show vlacp interface` command displays a column called `HAVE PARTNER` , with possible values of `yes` or `no` .

If `HAVE PARTNER` is `yes` when `ADMIN ENABLED` and `OPER ENABLED` are `true` , then that port has received VLACPDUs from a port and those PDUs were recognized as valid according to the interface settings.

If `HAVE PARTNER` is `no` , when `ADMIN ENABLED` and `OPER ENABLED` are **true**, then that port did not received any VLACPDUs yet.

If `HAVE PARTNER` is `no` , when `ADMIN ENABLED` is `true` and `OPER ENABLED` is `FALSE` , then the partner for that port is down (that port received at least one correct VLACPDU, but did not receive additional VLACPDUs within the configured timeout period). In this case VLACP blocks the port.

As long as the VLACP functional address for a specific interface is not changed when using the command (config-if)#vlacp port x funcmac-addr H.H.H/xx.xx.xx.xx.xx.xx, the MAC address is displayed as 00:00:00:00:00:00. The MAC address used for sending VLACP PDUs for an interface is the global VLACP MAC address (01:80:c2:00:11:00). The VLACP global destination MAC can be specified by the user. Setting a func-mac-addr on an interface displays that address in the show vlacp interface instead of 00:00:00:00:00:00.

Using Distributed Multi-Link Trunking

The following sections describe how to configure Distributed Multi-Link Trunking (DMLT) using the ACLI.

Displaying DMLT configuration

Use the following procedure to display Distributed Multi-Link Trunking (DMLT) configuration and utilization.

Procedure steps

To display DMLT configuration and utilization, use the following command from Global Configuration mode:

```
show mlt [utilization <1-6>] [spanning-tree <1-6>]
```

Variable definitions

Variable	Value
utilization <1-6>	Displays the utilization of the specified enabled MLT(s) in percentages.
spanning tree <1-6>	Displays Multi-Link trunk spanning tree settings.

Configuring DMLT

Use the following procedure to configure Distributed Multi-Link Trunking (DMLT).

Procedure steps

To configure DMLT, use the following command from Global Configuration mode:

```
mlt [<1-6> spanning-tree]
```

Variable definitions

Variable	Value
<1-6>	Specifies MLT ID
spanning tree	Sets MLT spanning-tree settings

Using Distributed Link Aggregation Group

The following sections contain procedures used to configure 802.3ad Link Aggregation using the ACLI.

Displaying LACP aggregator information

Use the following procedure to display LACP aggregator information.

Procedure steps

To display LACP aggregator information, use the following command from Global Configuration mode:

```
show lacp aggr [<1-65535>]
```

Variable definitions

Variable	Value
<1-65535>	Specifies the aggregator ID

Displaying LACP port debug information

Use the following procedure to display LACP port debug information.

Procedure steps

To display LACP port debug information, use the following command from Global Configuration mode:

```
show lacp debug member
```

Variable definitions

Variable	Value
portlist	Enter the ports for which you want debug information.

Displaying LACP port information

Use the following procedure to display LACP port information.

Procedure steps

To display LACP port information, use the following command from Global Configuration mode:

```
show lacp port
```

Variable definitions

Variable	Value
aggr	Selects port that are members of aggregator
portlist	Enter the ports for which you want information

Displaying LACP statistics information

Use the following procedure to display LACP statistics information.

Procedure steps

To display LACP statistics information, use the following command from Interface Configuration mode:

```
show lacp stats
```

Variable definitions

Variable	Value
aggr	Selects port that are members of aggregator
portlist	Enter the ports for which you want statistics

Displaying LACP system settings

Use the following procedure to display LACP system settings.

Procedure steps

To display LACP system settings, use the following command from Global Configuration mode:

```
show lacp system
```

Configuring LACP system priority

Use the following procedure to change the LACP system priority. Default is 32768.

Procedure steps

To configure LACP system priority, use the following command from Global Configuration mode:

```
lacp sys <0-65535>
```

Configuring the administrative key for a set of ports

Use the following procedure to configure the administrative key for a set of ports.

Procedure steps

To configure the administrative key, use the following command from Interface Configuration mode:

```
lacp key <port> <1-4095>
```

Variable definitions

Variable	Value
port	Specifies port list
<1-4095>	Key value

Configuring LACP priority

Use the following procedure to configure LACP priority for ports.

Procedure steps

To configure LACP priority, use the following command from Interface Configuration mode:

```
lacp priority <port> <0-65535>
```

Variable definitions

Variable	Value
port	Specifies port list
<0-65535>	Specifies LACP port priority

Configuring LACP operating mode

Use the following procedure to configure the LACP operating mode for a set of ports. Default is off.

Procedure steps

To configure the LACP operating mode, use the following command from Interface Configuration mode:

```
lacp mode [port <portlist>] {off | passive | active}
```

Variable definitions

Variable	Value
port	Enter the ports for which you want to set LACP the mode.
port {off | passive | active}	Sets the LACP mode for the specified port to off, passive, or active; if port mode is selected as Passive or Active, port is ready to participate in LACP

Configuring LACP timeout

Use the following procedure to configure the LACP timeout time for a set of ports. Default is Long.

Procedure steps

To configure LACP timeout, use the following command from Interface Configuration mode:

```
lacp timeout-time [port <portlist>] {short | long}
```

Variable definitions

Variable	Value
port	Enter the ports for which you want to set an LACP timeout.
{short | long}	Sets short or long LACP timeout for the port. The long timeout is 90 seconds and short timeout is 3 seconds.

Clearing LACP port statistics

The `lacp clear-stats` command clears port statistics.

Procedure steps

To clear LACP port statistics, use the following command from Interface Configuration mode:

```
lacp clear-stats [port <portlist>]
```

Chapter 11: Configuring ADAC for Avaya IP Phones using ACLI

You can configure ADAC-related settings using the ACLI. For more information about the ADAC feature, see ADAC Fundamentals on page 51

Configuring global ADAC settings

Use the following procedure to configure global ADAC settings for a device.

Procedure steps

To configure settings, use the following command from Global Configuration mode:

```
adac {[enable] [op-mode {untagged-frames-basic|untagged-
frames-advanced|tagged-frames}][voice-vlan <1-4094>][uplink-
port <portlist>][call-server-port <portlist>][mac-range-table
{low-end}{0123.4567.89ab}{high-end}{0123.4567.89zz}]}
```

Important:
MAC address must be entered in Hexadecimal format.

Variable definitions

Parameters and variables	Description
enable	Enables ADAC on the device.
op-mode {untagged-frames-basic|untagged-frames-advanced|tagged-frames}	Sets the ADAC operation mode to one of the following:

Parameters and variables	Description
	• untagged-frames-basic: IP Phones send untagged frames, and the Voice VLAN is not created • untagged-frames-advanced: IP Phones send untagged frames, and the Voice VLAN is created • tagged-frames: IP Phones send tagged frames, and the Voice VLAN is created
voice-vlan <1–4094>	Configures the Voice VLAN ID. The assigned VLAN ID must not previously exist.
uplink-port <portlist>	Configures a maximum of 8 ports as uplink ports.
call-server-port <portlist>	Configures a maximum of 8 ports as Call Server ports.
mac-range-table {low-end}{0123.4567.89ab} {high-end}{0123.4567.89zz}	Adds new supported MAC address range. **Important:** Specify the low-end parameter first to set the high-end parameter (H.H.H/ xx.xx.xx.xx.xx.xx) for mac-range-table.

Disabling or clearing ADAC settings

Use the following procedure to disable ADAC on the device or clear the ADAC settings for the device.

Procedure steps

To disable or clear settings, use the following command from Global Configuration mode:

```
no adac {[enable] [voice-vlan] [uplink-port] [call-server-port]
[mac-range-table {low-end}{0123.4567.89ab}{high-end}
{0123.4567.89zz}]}
```

Important:
MAC address must be entered in Hexadecimal format.

Variable definitions

Parameters and variables	Description
enable	Disables ADAC on the device
voice-vlan	Clears Voice-VLAN ID
uplink-port	Clears the uplink ports
call-server-port	Clears the Call Server ports
mac-range-table {low-end}{0123.4567.89ab} {high-end}{0123.4567.89zz}	Deletes the supported MAC address range **Important:** Specify the low-end parameter first to set the high-end parameter (H.H.H/ xx.xx.xx.xx.xx.xx) for mac-range-table.

Resetting ADAC settings to default

Use the following procedure to restore default ADAC settings on the device.

Procedure steps

To restore settings, use the following command from Global Configuration mode:

```
default adac {[enable][op-mode][voice-vlan][uplink-port][call-server-port][mac-range-table]
```

Variable definitions

Parameters and variables	Description
enable	Restores the default state of ADAC
op-mode	Restores the default ADAC operation mode

Parameters and variables	Description
voice-vlan	Restores the default Voice-VLAN ID
uplink-port	Restores the default Uplink port
call-server-port	Restores the default Call Server port
mac-range-table	Restores the MAC address ranges supported by default

Configuring ADAC MAC address ranges

Use the following procedure to add a specified range to the table of MAC addresses recognized asAvaya IP Phones by the Auto-Detection process.

Procedure steps

To add a range, use the following command from Global Configuration mode:

```
[no] adac mac-range-table low-end <aaaa.aaaa.aaaa> high-end
<bbbb.bbbb.bbbb>
```

Use the **no** form of this command to delete a range.

Variable definitions

Parameters and variables	Description
<low-end>	Specifies the low-end of the MAC address range
<high-end>	Specifies the high-end of the MAC address range

Resetting MAC address ranges

Use the following procedure to restore all supported MAC address ranges on the switch to their default values.

Procedure steps

To restore to default, use the following procedure from Global Configuration mode:

```
default adac mac-range-table
```

Configuring ADAC device settings per port

Use the following procedure to set ADAC settings for the device on a specific port.

Procedure steps

To configure settings, use the following command from Interface Configuration mode:

```
adac [port <portList>]{[enable][tagged-frames-pvid{<1-4094 >|
no-change}][tagged-frames-tagging{tagAll| tagPvidOnly|
untagPvidOnly| no-change}] [detection {[mac][lldp]}]}
```

Variable definitions

Parameters and variables	Description
enable	Enables auto-detection on ports
port <portlist>	Port number for which to change settings
tagged-frames-pvid {<1-4094 >|no-change}	Sets Tagged-Frames PVID on the port or ports listed Use no-change to keep the current setting
tagged-frames-tagging{tagAll|tagPvidOnly| untagPvidOnly|no-change}	Sets Tagged-Frames Tagging to • tagAll • tagPvidOnly • untagPvidOnly Use no-change to keep the current setting
detection {[mac][lldp]}	Enables detection mechanisms on ports • mac • lldp

Setting ADAC detection method

Use the following procedure to set the detection method, by MAC address or using LLDP (IEEE 802.1AB) for a device on a port.

Procedure steps

To set the detection method, use the following command from Interface Configuration mode:

```
[no] adac detection [port <portList>]{[mac][lldp]}
```

Use the **no** form of this command to disable.

Variable definitions

Parameters and variables	Description
lldp	Enables 802.1AB-based detection on ports
mac	Enables MAC-based detection on ports
port	Specifies the port or ports for which to set the detection mode

Disabling ADAC per port

Use the following procedure to disable ADAC settings for the device on a specific port.

Procedure steps

To disable ADAC settings, use the following procedure from Interface Configuration mode:

```
no adac [port <portList>][enable]}
```

Variable definitions

Parameters and variables	Description
port<portlist>	Specifies the port numbers for which to change the settings
enable	Disables auto-detection on ports

Resetting ADAC port settings to default

Use the following procedure to restore the per port ADAC settings to defaults for the specified ports.

Procedure steps

To restore settings, use the following command from Interface Configuration mode:

```
default adac [port <portList>]{[enable][tagged-frames-pvid]
[tagged-frames-tagging]}
```

Variable definitions

Parameters and variables	Description
port<portlist>	Specifies the port numbers for which to change the settings
enable	Restores default auto-detection on ports
tagged-frames-pvid	Restores default PVID to be configured for telephony ports in Tagged Frames operating mode
tagged-frames-tagging	Restores default tagging to be configured for telephony ports in Tagged Frames operating mode

Restoring ADAC detection method to default

Use the following procedure to restore the ADAC auto-detection method by either MAC address or LLDP for a device on a port.

Procedure steps

To restore the detection method, use the following command from Interface Configuration mode:

```
default adac detection [port <portlist>] {[mac] [lldp]}
```

Variable definitions

Parameters and variables	Description
lldp	Restores default 802.1AB-based detection on ports
mac	Restores default MAC-based detection on ports
port	Specifies the port numbers for which to change the settings

Displaying ADAC settings per port

Use the following procedure to display ADAC settings for the device on a specific port.

Procedure steps

To display settings, use the following command from Global Configuration mode:

```
show adac interface
```

Variable definitions

Parameters and variables	Description
Type	Specifies how ADAC classifies this port • T: Telephony port • CS: Call Server port • U: Uplink port or part of the same trunk as the current set uplink port
Auto-Detection	Controls whether the interface should auto-detect; if there is any Avaya IP Phone connected to it (and implicitly apply auto-configuration for it)
Oper State	Indicates whether ADAC is enabled or disabled on that port
Auto-Configuration	Specifies if the auto-configuration is applied on a port or not
Tagged-Frames PVID	Specifies the PVID value that Auto-Configuration apply for ports having Auto-Detection enabled and running in Tagged-Frames operational mode. A value of 0 indicates that Auto-Configuration cannot change the PVID for the respective port. If the VLAN with the ID equal with this PVID does not exist when Auto-Configuration is applied to a port, then Auto-Configuration won't change the port's PVID (it will ignore the current value of this parameter, and treat it as if its value is currently 0);
Tagged-FramesTagging	Specifies the tagging value that Auto-Configuration apply for ports having Auto-Detection enabled and running in Tagged-Frames operational mode.

Displaying ADAC MAC range

Use the following procedure to display the range of MAC addresses used by ADAC to identify an IP Phone with the MAC detection mechanism.

Procedure steps

To display the ranges, use the following command from Global Configuration mode:

```
show adac mac-range-table
```

Displaying ADAC detection method status

Use the following procedure to display the status of detection mechanism for the device on a specific port.

Procedure steps

To display status, use the following command from Interface Configuration mode:

```
show adac detection interface
```

ADAC and 802.1AB interoperability configuration example

The following sections provide examples of the steps required to properly configure ADAC and 802.1AB interoperability using ACLI.

Step 1: Configure LLDP MED Network Policy TLV transmission

Configure LLDP MED Network Policy TLV transmission on ports with IP phones (LLDP capable) connected:

```
2526T-PWR(config-if)# lldp tx-tlv med med-capabilities
2526T-PWR(config-if)# lldp tx-tlv med network-policy
2526T-PWR# show lldp port 5 tx-tlv med
-------------------------------------------------------------
                    lldp port med tlvs
-------------------------------------------------------------
-------------------------------------------------------------

Port     Med          Network   Location  Extended  Inventory
         Capabilitie  Policy              PSE
         s
```

```
5         true          true      false     false     false
```

MED TLVs are transmitted only if Med-Capabilities TLV is transmitted

Step 2: Configure ADAC

Configure ADAC globally and per port:

```
2526T-PWR(config)# adac voice-vlan 555

2526T-PWR(config) adac uplink-port 1

2526T-PWR(config)# adac op-mode tagged-frames

2526T-PWR(config)# adac enable

2526T-PWR(config)# show adac

                          ADAC Global
                         Configuration

------------------------------------------------------------

ADAC Admin State: Enabled
ADAC Oper State: Enabled
Operating Mode: Tagged Frames
Traps Control Status: Enabled
Voice-VLAN ID: 555
Call Server Port: None
Uplink Port: 1

Enable ADAC for telephony ports:

2526T-PWR(config-if)# adac enable

2526T-PWR(config)# show adac interface 5

Port   Type    Auto      Oper      Auto      T-F PVID  T-F
               Detectio  State     Config—             Tagging
               n                   uration

5      T       Enabled   Enabled   Applied   No        Untag
                                             Change    PVID
                                                       Only
```

Step 3: Enable Automatic QoS

Enable Automatic QoS mixed or pure mode:

```
2526T-PWR(config)# qos agent nt-mode mixed

2526T-PWR(config)# show qos agent
QoS Operational Mode: Enabled
QoS NVRam Commit Delay: 10 seconds
QoS Queue Set: 4
QoS Buffering: Maximum
QoS Default Statistics Tracking: Aggregate
QoS NT Mode: NT w/ egress DSCP remapping
QoS Trusted Processing Mode: Partial
```

Step 4: Verify the network policy

Verify the network policy sent for telephony ports (the DSCP is updated from 46 to 47):

```
2526T-PWR(config-if)show lldp local-sys-data med
-------------------------------------------------------------

               lldp local-sys-data chassis

 -------------------------------------------------------------
 -------------------------------------------------------------

      ChassisId: MAC address 00:1b:25:dd:58:00
      SysName:
      SysCap: rB / B (Supported/Enabled)
      SysDescr:

Ethernet Routing Switch    HW:01    FW:1.0.0.15
2526T-PWR                           SW:v4.3.0.032

     MED-Device class:        Network Connectivity Device

     MED-POE Device Type      PSE Device

     HWRev: 01                SerialNumber: LBNNTMJL4500CN

     FWRev: 1.0.0.15          SWRev: v4.3.0.032

     ManufName: Avaya         ModelName: 2526T-PWR

 -------------------------------------------------------------
```

```
                  lldp local-sys-data port
-------------------------------------------------------------

Port: 5

MED-Capabilities:          CNLSI

MED-PSE PDPort Priority:           Power Value: 16.0 Watt
Low

MED-Application Type:              VLAN ID: 555
Voice

L2 Priority: 6    DSCP Value: 47   Tagged Vlan, Policy
                                   defined

 -------------------------------------------------------------

Sys capability: O-Other; R-Repeater; B-Bridge; W-WLAN
accesspoint; r-Router; T-Telephone; D-DOCSIS cable device; S-
Station only.

Med Capabilities-C: N-Network Policy; L-Location
Identification; I-Inventory; S-Extended Power via MDI - PSE;
D-Extended Power via MDI - PD.
```

Chapter 12: VLAN configuration using Enterprise Device Manager

This chapter describes how to use Enterprise Device Manager (EDM) to manage VLANs on your Ethernet Routing Switch 2500 Series . This chapter covers creating, editing, and deleting VLANs.

Use Enterprise Device Manager to manage VLANs on your Ethernet Routing Switch 2500 Series . You can create, edit, and delete VLANs.

VLANs

A VLAN is a collection of ports on one or more switches that define a broadcast domain. The Ethernet Routing Switch 2500 Series supports port-based and IPv6 protocol-based VLANs.

When you create VLANs using Enterprise Device Manager, observe the following rules:

- The ports in a VLAN or Multi-Link trunk must be a subset of a Single Spanning Tree Group.
- VLANs must have unique VLAN IDs and names.

VLAN management using EDM

Use the information in this section to view, create, and manage VLAN configurations for a switch or stack.

Viewing VLAN information using EDM

Use this procedure to display VLAN configuration information for a switch or stack.

Procedure steps

1. From the navigation tree, choose **VLAN**.
2. Double-click **VLANs**.

3. In the work area, click the **Basic** tab.
4. To display IP address information for a VLAN, click the **VLAN ID**.
5. Click the **IP** button.
6. To display IPv6 address information for a VLAN, click the **VLAN ID**.
7. Click the **IPv6** button.

Variable definitions

Use the data in this table to help you understand the VLAN display.

Variable	Value
Id	Indicates the VLAN ID for the VLAN.
Name	Indicates the name of the VLAN.
IfIndex	Indicates the interface index. This is a read-only value.
Type	Indicates the type of VLAN. Values include: • byPort — VLAN by port • byProtocolId — VLAN by protocol ID
PortMembers	Indicates the ports that are members of the VLAN.
ActiveMembers	Indicates the ports that are currently active in the VLAN. Active ports include all static ports and any dynamic ports where the VLAN policy was met.
StgId	Indicates the STG to which the selected port(s) belong. **Important:** This column is available only when the switch is operating in the STPG mode. Ethernet Routing Switch 2500 Series does not support multiple STGs when operating in the STPG mode.
ProtocolId	Indicates the protocol identifier for the VLAN. The protocol ID is significant only when the VLAN type is byProtocolId, otherwise the protocol ID value is none (0). Values include: • 0 • ipV6
UserDefinedPid	Indicates the user defined protocol identifer for a protocol based VLAN.
MstpInstance	Indicates the MSTP instance associated with the VLAN. Values include: • none • cist • msti-1-7

Variable	Value
	Important: This column is available only when the switch is operating in the MSTP mode.
MacAddress	Indicates the MAC address associated with the VLAN.
Routing	Indicates whether routing is enabled (true) or disabled (false) for the VLAN.

Modifying an existing VLAN in STPG mode using EDM

Use this procedure to modify the configuration of an existing VLAN when the Spanning Tree administration operating mode is avayaStpg.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **VLANs**.
3. In the work area, click the **Basic** tab.
4. To select a VLAN to edit, click the **VLAN ID**.
5. In the VLAN row, double-click the cell in the **Name** column.
6. Type a character string to assign a unique name to the VLAN.
7. In the VLAN row, double-click the cell in the **PortMembers** column.
8. Select ports to add to the VLAN.

   OR

   Deselect ports to remove them from the VLAN.
9. Click **Ok** .
10. In the VLAN row, double-click the cell in the **Routing** column.
11. Select a value from the list—true to enable routing for the VLAN, or false to disable routing for the VLAN .
12. Click **Apply** .

Variable definitions

Use the data in this table to modify the configuration of an existing VLAN in STPG mode.

Variable	Value
Id	Indicates the ID for the VLAN. This is a read-only value.
Name	Specifies an alphanumeric name for the VLAN. If you do not type a name, the switch default is applied.
IfIndex	Indicates the interface index. This is a read-only value.
Type	Indicates the type of VLAN: byPort or byProtocolId. This is a read-only value.
PortMembers	Specifies the ports that are members of the VLAN.
ActiveMembers	Indicates the ports that are currently active in the VLAN. Active ports include all static ports and any dynamic ports where the VLAN policy was met. This is a read-only value.
StgId	Indicates the STG to which the selected port or ports belong. This is a read-only value. **Important:** This column is available only when the Spanning Tree administration operating mode is avayaStpg . The switch does not support multiple STGs when operating in the STPG mode. **Important:** When the Spanning Tree administration operating mode is RSTP, this column is not available.
MacAddress	Indicates the MAC address associated with the VLAN. This is a read-only value.
Routing	Specifies whether routing is enabled (true) or disabled (false) for the VLAN.

Modifying an existing VLAN in RSTP mode using EDM

Use this procedure to modify the configuration of an existing VLAN when the Spanning Tree administration operating mode is RSTP.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **VLANs**.
3. In the work area, click the **Basic** tab.
4. To select a VLAN to edit, click the VLAN ID.
5. In the VLAN row, double-click the cell in the **Name** column.

6. Type a character string to assign a unique name to the VLAN.
7. In the VLAN row, double-click the cell in the **PortMembers** column.
8. Select ports to add to the VLAN.

   OR

   Deselect ports to remove them from the VLAN.
9. Click **Ok** .
10. In the VLAN row, double-click the cell in the **Routing** column.
11. Select a value from the list—true to enable routing for the VLAN, or false to disable routing for the VLAN .
12. Click **Apply** .

Variable definitions

Use the data in this table to modify the configuration of an existing VLAN in RSTP mode.

Variable	Value
Id	Indicates the ID for the VLAN. This is a read-only value.
Name	Specifies an alphanumeric name for the VLAN. If you do not type a name, the switch default is applied.
IfIndex	Indicates the interface index. This is a read-only value.
Type	Indicates the type of VLAN: byPort or byProtocolId. This is a read-only value.
PortMembers	Specifies the ports that are members of the VLAN.
ActiveMembers	Indicates the ports that are currently active in the VLAN. Active ports include all static ports and any dynamic ports where the VLAN policy was met. This is a read-only value.
MacAddress	Indicates the MAC address associated with the VLAN. This is a read-only value.
Routing	Specifies whether routing is enabled (true) or disabled (false) for the VLAN.

Modifying an existing VLAN in MSTP mode using EDM

Use this procedure to modify the configuration of an existing VLAN when the Spanning Tree administration operating mode is MSTP.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **VLANs**.
3. In the work area, click the **Basic** tab.
4. To select a VLAN to edit, click the VLAN ID.
5. In the VLAN row, double-click the cell in the **Name** column.
6. Type a character string to assign a unique name to the VLAN.
7. In the VLAN row, double-click the cell in the **PortMembers** column.
8. Select ports to add to the VLAN.

   OR

   Deselect ports to remove them from the VLAN.
9. Click **Ok** .
10. In the VLAN row, double-click the cell in the **MstpInstance** column, if the switch is in MSTP mode.
11. Select a value from the list.
12. In the VLAN row, double-click the cell in the **Routing** column.
13. Select a value from the list—true to enable routing for the VLAN, or false to disable routing for the VLAN .
14. Click **Apply** .

Variable definitions

Use the data in this table to modify the configuration of an existing VLAN in MSTP mode.

Variable	Value
Id	Indicates the ID for the VLAN. This is a read-only value.
Name	Specifies an alphanumeric name for the VLAN. If you do not type a name, the switch default is applied.
IfIndex	Indicates the interface index. This is a read-only value.
Type	Indicates the type of VLAN: byPort or byProtocolId. This is a read-only value.
PortMembers	Specifies the ports that are members of the VLAN.

Variable	Value
ActiveMembers	Indicates the ports that are currently active in the VLAN. Active ports include all static ports and any dynamic ports where the VLAN policy was met. This is a read-only value.
MstpInstance	The MSTP instance associated with the VLAN. Values include: • none • cist • msti-1-7 **Important:** This column is available only when the Spanning Tree administration operating mode is MSTP. **Important:** When the Spanning Tree administration operating mode is RSTP, this column is not available.
MacAddress	Indicates the MAC address associated with the VLAN. This is a read-only value.
Routing	Specifies whether routing is enabled (true) or disabled (false) for the VLAN.

Creating a VLAN in STP mode using EDM

Use the following procedure to create a new VLAN when the switch is in STP mode.

Prerequisites

Select avayaStpg for the Spanning Tree administration mode.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **VLANs**.
3. In the work area, click the **Basic** tab.
4. Click **Insert**.
5. In the **Id** field, type a value.

   OR

   Accept the default ID for the VLAN.
6. In the **Name** field, type a value.

OR

Accept the default name for the VLAN.

7. In the **Type** field, select **byPort** or **byProtocolId**.
8. Click **Insert**.
9. In the VLAN row, double-click the cell in the **PortMembers** column.
10. Select ports to add to the VLAN.

    OR

    Deselect ports to remove them from the VLAN.
11. Click **Ok** .
12. In the VLAN row, double-click the cell in the **Routing** column.
13. Select a value from the list—**true** to enable routing for the VLAN, or **false** to disable routing for the VLAN .
14. Click **Apply** .

Variable definitions

Use the data in this table to modify the create VLAN in STPG mode.

Variable	Value
Id	Specifies the ID for the VLAN.
Name	Specifies an alphanumeric name for the VLAN. If you do not type a name, the switch default is applied.
IfIndex	Indicates the interface index. This is a read-only value.
Type	Indicates the protocol identifier for the VLAN. The protocol ID is significant only when the VLAN type is byProtocolId. The only supported value is ipv6.
PortMembers	Specifies the ports that are members of the VLAN.
ActiveMembers	Indicates the ports that are currently active in the VLAN. Active ports include all static ports and any dynamic ports where the VLAN policy was met. This is a read-only value.
StgId	Indicates the STG to which the selected port or ports belong. This is a read-only value. **Important:** This column is available only when the Spanning Tree administration operating mode is avayaStpg . The switch does not support multiple STGs when operating in the STPG mode.

Variable	Value
	Important: When the Spanning Tree administration operating mode is RSTP, this column is not available.
ProtocolId	Indicates the protocol identifier for the VLAN. The protocol ID is significant only when the VLAN type is byProtocolId, otherwise the protocol ID value is none (0). Values include: • 0 • ipV6
UserDefinedPid	Indicates the user defined protocol identifier for a protocol based VLAN.
MacAddress	Indicates the MAC address associated with the VLAN. This is a read-only value.
Routing	Specifies whether routing is enabled (true) or disabled (false) for the VLAN.

Creating a VLAN in RSTP mode using EDM

Use the following procedure to create a new VLAN when the switch is in RSTP mode.

Prerequisites

Select avayaStpg for the Spanning Tree administration mode.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **VLANs**.
3. In the work area, click the **Basic** tab.
4. Click **Insert**.
5. In the **Id** dialog box, type a value.

   OR

   Accept the default ID for the VLAN.
6. In the **Name** dialog box, type a value.

OR

Accept the default name for the VLAN.

7. In the **Type** field, select **byPort** or **byProtocolId**.
8. Click **Insert**.
9. In the VLAN row, double-click the cell in the **PortMembers** column.
10. Select ports to add to the VLAN.

    OR

    Deselect ports to remove them from the VLAN.
11. Click **Ok** .
12. In the VLAN row, double-click the cell in the **Routing** column.
13. Select a value from the list—**true** to enable routing for the VLAN, or **false** to disable routing for the VLAN .
14. Click **Apply** .

Variable definitions

Use the data in this table to modify the create a VLAN in RSTP mode.

Variable	Value
Id	Specifies the ID for the VLAN.
Name	Specifies an alphanumeric name for the VLAN. If you do not type a name, the switch default is applied.
IfIndex	Indicates the interface index. This is a read-only value.
Type	Indicates the protocol identifier for the VLAN. The protocol ID is significant only when the VLAN type is byProtocolId. The only supported value is ipv6.
PortMembers	Specifies the ports that are members of the VLAN.
ActiveMembers	Indicates the ports that are currently active in the VLAN. Active ports include all static ports and any dynamic ports where the VLAN policy was met. This is a read-only value.
ProtocolId	Indicates the protocol identifier for the VLAN. The protocol ID is significant only when the VLAN type is byProtocolId, otherwise the protocol ID value is none (0). Values include: • 0 • ipV6
UserDefinedPid	Indicates the user defined protocol identifier for a protocol based VLAN.

Variable	Value
MacAddress	Indicates the MAC address associated with the VLAN. This is a read-only value.
Routing	Specifies whether routing is enabled (true) or disabled (false) for the VLAN.

Creating a VLAN in MSTP mode using EDM

Use the following procedure to create a new VLAN when the switch is in MSTP mode.

Prerequisites

Select MSTP for the Spanning Tree administration mode.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **VLANs**.
3. In the work area, click the **Basic** tab.
4. Click **Insert**.
5. In the **Id** dialog box, type a value.

   OR

   Accept the default ID for the VLAN.
6. In the **Name** dialog box, type a value.

   OR

   Accept the default name for the VLAN.
7. In the **Type** field, select **byPort** or **byProtocolId**.
8. Click the **MstpInstance** box arrow.
9. Select a value from the list.
10. Click **Insert**.
11. In the VLAN row, double-click the cell in the **PortMembers** column.
12. Select ports to add to the VLAN.

    OR

Deselect ports to remove them from the VLAN.

13. Click **Ok** .
14. In the VLAN row, double-click the cell in the **Routing** column.
15. Select a value from the list—**true** to enable routing for the VLAN, or **false** to disable routing for the VLAN .
16. Click **Apply** .

Variable definitions

Use the data in this table to modify the create a VLAN in MSTP mode.

Variable	Value
Id	Specifies the ID for the VLAN.
Name	Specifies an alphanumeric name for the VLAN. If you do not type a name, the switch default is applied.
IfIndex	Indicates the interface index. This is a read-only value.
Type	Indicates the protocol identifier for the VLAN. The protocol ID is significant only when the VLAN type is byProtocolId. The only supported value is ipv6.
PortMembers	Specifies the ports that are members of the VLAN.
ActiveMembers	Indicates the ports that are currently active in the VLAN. Active ports include all static ports and any dynamic ports where the VLAN policy was met. This is a read-only value.
ProtocolId	Indicates the protocol identifier for the VLAN. The protocol ID is significant only when the VLAN type is byProtocolId, otherwise the protocol ID value is none (0). Values include: • 0 • ipV6
UserDefinedPid	Indicates the user defined protocol identifier for a protocol based VLAN.
MstpInstance	The MSTP instance associated with the VLAN. Values include: • none • cist • msti-1-7 **Important:** This column is available only when the Spanning Tree administration operating mode is MSTP. **Important:** When the Spanning Tree administration operating mode is RSTP, this column is not available.

Variable	Value
MacAddress	Indicates the MAC address associated with the VLAN. This is a read-only value.
Routing	Specifies whether routing is enabled (true) or disabled (false) for the VLAN.

Deleting a VLAN using EDM

Use this procedure to delete a VLAN.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **VLANs**.
3. To select a VLAN to delete, click the VLAN ID.
4. Click **Delete**.
5. Click **Yes**.

VLAN configuration for ports using EDM

Use the information in this section to view and configure VLAN membership for specific ports.

Viewing VLAN membership port information using EDM

Use this procedure to display the VLAN membership information for switch ports.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **VLANs** .
3. Click the **Ports** tab.

Variable definitions

Use the data in the following table to help you understand the VLAN port membership display.

Variable	Value
Index	Indicates the switch position in the stack and the port number. This is a read-only value.
VlanIds	Indicates the VLAN IDs of which this port is a member. This is a read-only value.
DiscardUntaggedFrames	Indicates how untagged frames received on this port are processed. • true—untagged frames are discarded by the forwarding process • false—untagged frames are assigned to the VLAN specified by the VLAN ID. This column applies to trunk ports only.
FilteredUnregisteredFrame	Indicates how unregistered frames received on this port are processed. • true—unregistered frames are discarded by the forwarding process • false—unregistered frames are assigned to the VLAN specified by the VLAN ID. This column applies to access ports only.
DefaultVlanId	Indicates the VLAN ID assigned to untagged and unregistered frames received on a port.
PortPriority	Indicates the port priority for the switch to consider as it forwards received packets. Values range from 0 to 7.
Tagging	Indicates the type of VLAN port. Possible values are: • untagAll (access) • tagAll (trunk) • untagPvidOnly • tagPvidOnly If the port is a trunk port, the port is often a member of more than one VLAN. If the port is an access port, the port can only be a member of one VLAN.

Configuring VLAN membership ports using EDM

Use this procedure to configure VLAN membership for one or more switch ports.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **VLANs** .
3. Click the **Ports** tab.
4. To select a port to edit, click the port row.
5. In the port row, double-click the cell in the **DiscardUntaggedFrames** column.
6. Select a value from the list—true to discard untagged frames for the port, or false to accept untagged frames for the port.
7. In the port row, double-click the cell in the **FilteredUnregisteredFrame** column.
8. Select a value from the list—true to discard unregistered frames for the port, or false to process unregistered frames normally for the port.
9. In the port row, double-click the cell in the **DefaultVlanId** column.
10. Type a value for the default VLAN ID.
11. In the port row, double-click the cell in the **PortPriority** column.
12. Select a value from the list.
13. In the port row, double-click the cell in the **Tagging** column.
14. Select a value from the list.
15. Repeat steps 5 through 15 to configure VLAN memberships for additional ports.
16. Click **Apply** .

Variable definitions

Variable	Value
Index	Indicates the switch position in the stack and the port number. This is a read-only value.
VlanIds	Indicates the VLAN IDs of which this port is a member. This is a read-only value.
DiscardUntaggedFrames	Specifies how untagged frames received on this port are processed.

Variable	Value
	• true—untagged frames are discarded by the forwarding process • false—untagged frames are assigned to the VLAN specified by the VLAN ID. This column applies to trunk ports only.
FilteredUnregisteredFrame	Specifies how unregistered frames received on this port are processed. • true—unregistered frames are discarded by the forwarding process • false—unregistered frames are assigned to the VLAN specified by the VLAN ID. This column applies to access ports only.
DefaultVlanId	Specifies the VLAN ID assigned to untagged and unregistered frames received on a port.
PortPriority	Specifies the port priority for the switch to consider as it forwards received packets. Values range from 0 to 7.
Tagging	Specifies the type of VLAN port. Possible values are: • untagAll (access) • tagAll (trunk) • untagPvidOnly • tagPvidOnly If the port is a trunk port, the port is often a member of more than one VLAN. If the port is an access port, the port can only be a member of one VLAN.

Selecting VLAN configuration control using EDM

Use the following procedure to select configuration control for a VLAN.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **VLANs**.
3. In the work area, click the **Settings** tab.

4. In the **ManagementVlanID** dialog box, type a value.
5. In the **VlanConfigControl** section, click a radio button.
6. Click **Apply** .

Variable definitions

Use the data in this table to set VLAN configuration control.

Variable	Value
ManagementVlanID	Specifies the identifier of the management VLAN. Values range from 1 to 4094.
VlanConfigControl	VlanConfigControl presents four selections: • automatic: This selection automatically adds an untagged port to a new VLAN and automatically removes it from any previous VLAN membership. The PVID of the port is automatically changed to the new VID of the VLAN it joins. Since the port is first added to the new VLAN and then removed from any previous membership, the Spanning Tree Group participation of the port is not disabled as long as the VLANs involved are in the same Spanning Tree Group • autopvid: When an untagged port is added to a new VLAN, the port is added to the new VLAN and the PVID is assigned to the new VID without removing it from any previous VLAN memberships. Using this option, an untagged port can have membership in multiple VLANs. • flexible: This selection functions in a similar manner to disabling AutoPVID functionality. When this option is used, an untagged port can belong to an unlimited number of VLANs. Any new additions of an untagged port to a new VLAN do not change the PVID of that port. • strict: The factory default, this selection restricts the addition of an untagged port to a VLAN if it is already a member of another VLAN. To add an untagged port to a new VLAN, the switch administrator must remove the port from all other VLANS of which it is a member before adding it to a new VLAN. The PVID of the port is changed to the new VID to which it was added.

Port configuration for VLANs using EDM

Use the information in this section to view and configure specific ports for VLAN membership.

Viewing port VLAN membership information using EDM

Use this procedure to display the VLAN membership information for switch ports.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis** .
3. In the Chassis tree, double-click **Ports** .
4. Click the **VLAN** tab.

Variable definitions

Variable	Value
Index	Indicates the switch position in the stack and the port number. This is a read-only value.
VlanIds	Indicates the VLAN IDs of which this port is a member. This is a read-only value.
DiscardUntaggedFrames	Indicates how untagged frames received on this port are processed. • true—untagged frames are discarded by the forwarding process • false—untagged frames are assigned to the VLAN specified by the VLAN ID. This column applies to trunk ports only.
FilteredUnregisteredFrame	Indicates how unregistered frames received on this port are processed.

Variable	Value
	• true—unregistered frames are discarded by the forwarding process • false—unregistered frames are assigned to the VLAN specified by the VLAN ID. This column applies to access ports only.
DefaultVlanId	Indicates the VLAN ID assigned to untagged and unregistered frames received on a port.
PortPriority	Indicates the port priority for the switch to consider as it forwards received packets. Values range from 0 to 7.
Tagging	Indicates the type of VLAN port. Possible values are: • untagAll (access) • tagAll (trunk) • untagPvidOnly • tagPvidOnly If the port is a trunk port, the port is often a member of more than one VLAN. If the port is an access port, the port can only be a member of one VLAN.

Configuring ports for VLAN membership using EDM

Use this procedure to configure one or more switch ports for VLAN membership.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis** .
3. In the Chassis tree, double-click **Ports** .
4. Click the **VLAN** tab.
5. To select a port to edit, click the port row.
6. In the port row, double-click the cell in the **DiscardUntaggedFrames** column.
7. Select a value from the list—true to discard untagged frames for the port, or false to accept untagged frames for the port.
8. In the port row, double-click the cell in the **FilteredUnregisteredFrame** column.
9. Select a value from the list—true to discard unregistered frames for the port, or false to process unregistered frames normally for the port.

10. In the port row, double-click the cell in the **DefaultVlanId** column.
11. Type a value for the default VLAN ID.
12. In the port row, double-click the cell in the **PortPriority** column.
13. Select a value from the list.
14. In the port row, double-click the cell in the **Tagging** column.
15. Select a value from the list.
16. Repeat steps 5 through 15 to configure VLAN memberships for additional ports.
17. Click **Apply** .

Variable definitions

Use the data in the following table to configure ports for VLAN membership.

Variable	Value
Index	Indicates the switch position in the stack and the port number. This is a read-only value.
VlanIds	Indicates the VLAN IDs of which this port is a member. This is a read-only value.
DiscardUntaggedFrames	Specifies how untagged frames received on this port are processed. • true—untagged frames are discarded by the forwarding process • false—untagged frames are assigned to the VLAN specified by the VLAN ID. This column applies to trunk ports only.
FilteredUnregisteredFrame	Specifies how unregistered frames received on this port are processed. • true—unregistered frames are discarded by the forwarding process • false—unregistered frames are assigned to the VLAN specified by the VLAN ID. This column applies to access ports only.
DefaultVlanId	Specifies the VLAN ID assigned to untagged and unregistered frames received on a port.
PortPriority	Specifies the port priority for the switch to consider as it forwards received packets. Values range from 0 to 7.
Tagging	Specifies the type of VLAN port. Possible values are:

Variable	Value
	• untagAll (access) • tagAll (trunk) • untagPvidOnly • tagPvidOnly If the port is a trunk port, the port is often a member of more than one VLAN. If the port is an access port, the port can only be a member of one VLAN.

MAC address table management using EDM

This section describes how to manage the MAC address table by clearing entries.

Important:

In certain situations, due to the hash algorithm used by switch to store MAC addresses into memory, some MAC addresses can not be learned.

Flushing the MAC address table using EDM

Use the following procedure to flush the MAC address table to clear all addresses in the MAC address table.

Procedure steps

1. From the navigation tree, click **Edit** to open the Edit navigation tree.
2. Double-click **Bridge** to open the Bridge work area.
3. Select the **MAC Flush** tab.
4. To clear all MAC address table entries, select the **FlushMacAddrTableAll** check box.
5. Click **Apply**.

Flushing the MAC address table for a FastEthernet interface using EDM

Use the following procedure to flush the MAC address table for a FastEthernet interface to clear the MAC address table for specified interface ports.

Procedure steps

1. From the navigation tree, click **Edit** to open the Edit navigation tree.
2. Double-click **Bridge** to open the Bridge work area.
3. Select the **MAC Flush** tab.
4. Click the **FlushMacAddrTableByPortList** elipsis **(...)**.
5. Select interface ports for which to clear MAC address table entries.
6. Click **Ok**.
7. Click **Apply**.

Flushing the MAC address table for a VLAN using EDM

Use the following procedure to flush the MAC address table for a VLAN to clear all MAC addresses for a specific VLAN.

Procedure steps

1. From the navigation tree, click **Edit** to open the Edit navigation tree.
2. Double-click **Bridge** to open the Bridge work area.
3. Select the **MAC Flush** tab.
4. Type a VLAN ID for which to clear the MAC address table in the **FlushMacAddrTableByVlan** box.
5. Click **Apply**.

Variable definitions

Use the data in this table to help you flush the MAC address table for a VLAN.

Variable	Value
FlushMacAddrTableByVlan	Specifies the VLAN ID. Values range from 1 to 4094.

Flushing the MAC address table for a trunk using EDM

Use the following procedure to flush the MAC address table for a trunk to clear all MAC addresses for members of a multi-link trunk.

Procedure steps

1. From the navigation tree, click **Edit** to open the Edit navigation tree.
2. Double-click **Bridge** to open the Bridge work area.
3. Select the **MAC Flush** tab.
4. Type a trunk number for which to clear the MAC address table in the **FlushMacAddrTableByTrunk** box.
5. Click **Apply**.

Variable definitions

Use the data in this table to help you flush the MAC address table for a VLAN.

Variable	Value
FlushMacAddrTableByTrunk	Specifies the multi-link trunk. Values range from 1 to 6.

Flushing a single MAC address table entry using EDM

Use the following procedure to flush a single MAC address table entry to clear one MAC address from the MAC address table.

Procedure steps

1. From the navigation tree, click **Edit** to open the Edit navigation tree.
2. Double-click **Bridge** to open the Bridge work area.
3. Select the **MAC Flush** tab.

4. Type a MAC address in the **FlushMacAddrTableByAddress** box.
5. Click **Apply**.

Variable definitions

Use the data in this table to help you flush the MAC address table for a VLAN.

Variable	Value
FlushMacAddrTableByAddress	Specifies a MAC address. The default value is 00:00:00:00:00:00.

Chapter 13: STP configuration using Enterprise Device Manager

This chapter describes using Enterprise Device Manager (EDM) to manage Spanning Tree Groups (STGs) on your Ethernet Routing Switch 2500 Series . It also discusses Rapid Spanning Tree protocol (RSTP), and the Multiple Spanning Tree Protocol (MSTP).

Changing the Spanning Tree mode using EDM

You can use the Enterprise Device Manager to change the Spanning Tree mode for the Ethernet Routing Switch 2500 Series.

Use the following procedure to change Spanning Tree mode.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **Spanning Tree**.
3. In the Spanning Tree navigation tree, double-click **Globals**.
4. In the **SpanningTreeAdminMode** section, click a radio button.
5. Click **Apply**.

   A warning message appears reminding you that you must reset the switch for the change to take effect.
6. Click **Yes**.
7. Reset the switch.

   For information about how to reset the switch, see Resetting the switch using EDM on page 158.
8. Rediscover the switch.

   For information about how to rediscover the switch, see Rediscovering the switch using EDM on page 158.

Resetting the switch using EDM

Use the following procedure to reset the switch.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Chassis**.
4. In the work area, click the **System** tab.
5. In the ReBoot section, click the **reboot** radio button.
6. Click **Apply**.

Note:
The rebooting process can take several minutes.

Rediscovering the switch using EDM

Use the following procedure to rediscover the switch after performing the switch reset procedure.

Procedure steps

1. From the navigation tree, double-click **Device**.
2. Double-click **Rediscover Device**.

Note:
The rediscover process can take several minutes.

Configuring STP BPDU Filtering using EDM

You can use the STP BPDU-Filtering tab to configure STP BPDU Filtering on a port. This tab is available in all three STP modes.

Use the following procedure to configure STP BPDU Filtering.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Ports**.
4. On the work area, click the **STP BPDU-Filtering** tab.
5. In the table, double-click a cell under the column heading for the parameter you want to change.
6. Select a parameter or value from the list.
7. Repeat the previous 2 steps until you have amended all of the parameters you want to change.
8. On the toolbar, click **Apply**.

Variable definitions

Variable	Value
rcPortIndex	Indicates the switch and port number.
AdminEnabled	Enables and disables BPDU filtering on the port.
OperEnabled	Indicates the current operational status of BPDU filtering on the port: true (enabled) or false (disabled).
Timeout	When BPDU filtering is enabled, this indicates the time (in 1/100 seconds) during which the port remains disabled after it receives a BPDU. The port timer is disabled if this value is set to 0. The default value is 12000 (120 seconds).
TimerCount	Displays the time remaining for the port to stay in the disabled state after receiving a BPDU.

Spanning Tree Group configuration using EDM

Use the information in this section to configure and manage a Spanning Tree Group (STG).

Configuring STG globally using EDM

Use the following procedure to configure Spanning Tree Group (STG) globally to select the STG configuration for the switch.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **STG** to open the STG work area.
4. Select the **Globals** tab.
5. Select a **SpanningTreePathCostCalculationMode** radio button.
6. Select a **SpanningTreePortMode** radio button.
7. Select or clear the **port802dot1dLearning** check box as required.
8. Click **Apply**.

Variable definitions

Variable	Value
SpanningTreePathCostCalculation Mode	This object indicates the current spanning-tree path cost calculation mode. Values include: • ieee802dot1dCompatible • ieee802dot1tCompatible The value ieee802dot1dCompatible is valid only after the switch is running in Avaya STPG mode.
SpanningTreePortMode	Specifies the STP port mode. Values include: • normal • auto

Variable	Value
SpanningTreeAdminCompatibility	Specifies the STP compatibility mode for various features. If port802dot1dLearning is selected, the port is goes to a Disabled state when the port operational status fails. If port802dot1dLearning is not selected, the port is remains in the Forwarding state when the port operational status fails.
SpanningTreeOperCompatibility	Indicates the STP compatibility mode for various features.

STG configuration tab

The Configuration tab in the STG dialog box contains general information for the STG.

Use the following procedure to view the **Configuration** tab.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **STG** to open the STG work area.
4. Select the **Configuration** tab.

Variable definitions

Variable	Value
Id	An identifier used to identify an STG in the device.
BridgeAddress	MAC address used by a bridge when it is referred to in a unique fashion. Avaya recommends that the number has to be the smallest MAC address of all ports belonging to the bridge. However, it is only required to be unique. When concatenated with Priority, a unique bridge identifier is formed that is used in the Spanning Tree Protocol.
NumPorts	Number of ports controlled by this bridging entity.
Protocol Specification	Version of the spanning tree protocol being run. Values include: • decLb100: Indicates the DEC LANbridge 100 Spanning Tree Protocol. • ieee8021d: IEEE 802.1d implementations will return this entry. When future versions of the IEEE Spanning Tree Protocol are

Variable	Value
	released that are incompatible with the current version, a new value will be defined.
Priority	Value of the writable portion of the bridge ID. That is, the first two octets of the (8-octet long) bridge ID. The last six octets of the bridge ID are given by the value of BridgeAddress.
BridgeMaxAge	Value, in units of hundredths of a second, that all bridges use for the maximum age of a bridge when it is acting as the root. **Important:** 802.1D-1990 specifies that the range is related to the value of BridgeHelloTime. The granularity of this timer is specified by 802.1D-1990 to be 1 second. A badValue error can be returned if the value set is not a whole number.
BridgeHelloTime	Value, in units of hundredths of a second, that all bridges use for HelloTime when a bridge is acting as the root. **Important:** The granularity of this timer is specified by 802.1D-1990 to be 1 second. A badValue error can be returned if the value set is not a whole number.
BridgeForward Delay	Value, in units of hundredths of a second, that all bridges use for ForwardDelay when this bridge is acting as the root. **Important:** 802.1D-1990 specifies that the range is related to the value of BridgeHelloTime. The granularity of this timer is specified by 802.1D-1990 to be 1 second. A badValue error can be returned if the value set is not a whole number.

STG status tab

The **Status** tab in the STG dialog box has status information for the STG.

Use the following procedure to view the **Status** tab.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **STG** to open the STG work area.
4. Select the **Status** tab.

Variable definitions

Variable	Value
Id	An identifier used to identify an STG in the device.
BridgeAddress	MAC address used by a bridge when it is referred to in a unique fashion. Avaya recommends that the number be the smallest MAC address of all ports belonging to the bridge. However, it is only required to be unique. When concatenated with Priority, a unique bridge identifier is formed that is used in the Spanning Tree Protocol.
NumPorts	Number of ports controlled by this bridging entity.
Protocol Specification	Version of the Spanning Tree Protocol being run. Values include: • decLb100: Indicates the DEC LANbridge 100 spanning tree protocol. • ieee8021d: IEEE 802.1d implementations will return this entry. When future versions of the IEEE spanning tree protocol are released that are incompatible with the current version, a new value will be defined.
TimeSinceTopology Change	Time (in hundredths of seconds) since the last topology change was detected by the bridge entity.
TopChanges	Number of topology changes detected by the bridge since the management entity was last reset or initialized.
DesignatedRoot	Bridge identifier of the root of the spanning tree as determined by the Spanning Tree Protocol. The value is used as the root identifier parameter in all configuration bridge PDUs originated by this node.
RootCost	Cost of the path to the root as seen from the bridge.
RootPort	Port that has the lowest cost path from the bridge to the root bridge.
MaxAge	Maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded, in units of hundredths of a second. This is the actual value that this bridge is currently using.
HelloTime	Amount of time between the transmission of configuration bridge PDUs by this node on any port when it is the root of the spanning tree (in hundredths of a seconds). This is the actual value that this bridge is currently using.

Variable	Value
HoldTime	Value of the interval length during which no more than two configuration bridge PDUs shall be transmitted by this node (in hundredths of a second).
ForwardDelay	This time value (in hundredths of a seconds) that controls how fast a port changes its spanning state when moving towards the forwarding state. Value determines how long the port stays in each of the listening and learning states, which precede the forwarding state. This is also used when a topology change has been detected and is underway, to age all dynamic entries in the forwarding database. **Important:** This value is the one that this bridge is currently using, in contrast to BridgeForwardDelay which is the value that this bridge and all others would start using if/when this bridge were to become the root.

STG Ports tab

The **Ports** tab in the STG dialog box has port information for the STG.

Use the following procedure to view the **Ports** tab.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **STG** to open the STG work area.
4. Select the **Ports** tab.

Variable definitions

Variable	Value
Port	Indicates the switch position in a stack and the port number. For a standalone switch, the default value of 1 is used for the switch position.
Stgld	STG identifier assigned to this port.

Variable	Value
Priority	Value of the priority field contained in the first octet of the port ID. The other octet is given by the value of the "rcStgPort."
State	The current state of the port as defined by application of the Spanning Tree Protocol. These are the instructions the port takes on a frame when it is received. If the bridge detects a port is malfunctioning, it will list it as "broken(6)." For ports that are disabled, the value is "disabled(1)."
EnableStp	Enables (True) or disables (False) the spanning tree of the port.
FastStart	When this is enabled (True), the port is move to forwarding or blocking state in 4 seconds.
AdminPathCost	The administrative value of PathCost.
PathCost	Contribution of the port to the pathcost of paths towards the spanning tree root, including the current port. 802.1D-1990 specifications recommends that the default of this parameter be in inverse proportion to the speed of the attached LAN.
DesignatedRoot	The unique "Bridge Identifier." This is recorded as Root in the configuration bridge PDUs transmitted by the Designated Bridge for the segment to that the port is attached.
DesignatedCost	Path cost of the Designated Port of the segment connected to the port. The value is compared to the Root Path Cost field in received bridge PDUs.
DesignatedBridge	Bridge identifier of the bridge that this port considers to be the Designated Bridge for this port's segment.
DesignatedPort	Port identifier of the port on the Designated Bridge for this port's segment.
Forward Transitions	Number of times this port has transitioned from the learning state to the forwarding state.

Configuring STG for a single port using EDM

Important:

You can access the **STG** tab only after the switch is operating in the STG mode.

In the **STG** tab, you can view the status and modify the configuration of a port's spanning tree parameters.

Use the following procedure to view the **STG** tab.

Procedure steps

1. From the Device Physical View, right click a port.
2. Click **Edit**.
3. In the Edit tree, click **Chassis**.
4. In the Chassis tree, click **Ports**.
5. In the work area, click the **STG** tab.
6. To select an STG to edit, click the STG ID.
7. In the STG row, double-click the cell in the **Priority** column.
8. Type a priority value.
9. In the STG row, double-click the cell in the **EnableStp** column.
10. Select a value from the list—true to enable STP for the STG, or false to disable STP for the STG.
11. In the STG row, double-click the cell in the **FastStart** column.
12. Select a value from the list—true to enable fast start for the STG, or false to disable fast start for the STG.
13. In the STG row, double-click the cell in the **AdminPathCost** column.
14. Type an administrative path cost value.
15. In the STG row, double-click the cell in the **PathCost** column.
16. Type a path cost value.
17. Click **Apply** .

Variable definitions

Variable	Value
StgId	Indicates the STG identifier assigned to this port. This is a read-only value.
Priority	Value of the priority field contained in the first octet of the port ID. The other octet is given by the value of the "rcStgPort."
State	Indicates the current port state as defined by application of the Spanning Tree Protocol. This state controls the action a port takes after it receives a frame. If the bridge detects a port that is malfunctioning, it places that port into the broken state. For ports that are disabled (see EnableStp), this object has a value of disabled. This is a read-only value.

Variable	Value
EnableStp	Enables (true) or disables (false) STP for the port.
FastStart	Enables (true) or disables (false) fast start for the port.
AdminPathCost	Specifies the administrative value of PathCost.
PathCost	Specifies the contribution of this port to the cost of paths toward the spanning tree root, which include this port. The IEEE 802.1D-1990 standard recommends that the default value of this parameter be in inverse proportion to the speed of the attached LAN.
DesignatedRoot	The unique Bridge Identifier of the bridge recorded as the Root in the Configuration BPDUs transmitted by the Designated Bridge for the segment to which the port is attached. This is a read-only value.
DesignatedCost	The path cost of the Designated Port of the segment connected to this port. This value is compared to the Root Path Cost field in received bridge PDUs. This is a read-only value.
DesignatedBridge	The Bridge Identifier of the bridge that this port considers to be the Designated Bridge for this port's segment. This is a read-only value.
DesignatedPort	The Port Identifier of the port on the Designated Bridge for this port's segment. This is a read-only value.
ForwardTransitions	The number of times this port has transitioned from the Learning state to the Forwarding state. This is a read-only value.

Rapid Spanning Tree Protocol

The Rapid Spanning Tree Protocol (RSTP or IEEE 802.1w) reduces the recovery time after a network break down. It also maintains a backward compatibility with the IEEE 802.1d which was the Spanning Tree implementation prior to RSTP. In certain configurations the recovery time of RSTP can be reduced to less than 1 second.

RSTP also reduces the amount of flooding in the network by enhancing the way Topology Change Notification (TCN) packet is generated.

Important:

You can access the RSTP menu command only after the switch is operating in the RSTP mode.

RSTP Globals tab

The **Globals** tab in the **RSTP** dialog box provides general information about RSTP when RSTP is the active mode.

Use the following procedure to view the **Globals** tab.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **RSTP**.

Variable definitions

Variable	Value
PathCost Default	Sets the version of the Spanning Tree default Path Costs that the Bridge uses. The value of 16-bit uses the 16-bit default Path Costs from IEEE Std. 802.1D-1998. A value of 32-bit uses the 32-bit default Path Costs from IEEE Std. 802.1t.
TxHoldCount	The value used by the Port Transmit state machine to limit the maximum transmission rate. The value can be between 1 to 10.
Version	The version of the Spanning Tree Protocol the bridge is currently running. The value 'stpCompatible' indicates that the bridge uses the Spanning Tree Protocol specified in IEEE 802.1D. The value 'rstp' indicates that the bridge uses Rapid Spanning Tree Protocol specified in IEEE 802.1w.
Priority	The value of the writable portion of the Bridge Identifier comprising of the first two octets. The values that are set for Priority must be in steps of 4096.
BridgeMaxAge	The value that all bridges use for MaxAge when this bridge is acting as the root. The granularity of this timer is specified to be 1 second. An agent can return a badValue error if a set is attempted to a value which is not a whole number of seconds
BridgeHelloTime	The value that all bridges use for HelloTime when this bridge is acting as the root. The granularity of this timer is specified

Variable	Value
	by 802.1D-1990 to be 1 second. An agent can return a badValue error if a set is attempted to a value which is not a whole number of seconds. Reference IEEE 802.1D-1990: Section 4.5.3.9.
BridgeForward Delay	The value that all bridges use for ForwardDelay when this bridge is acting as the root. Note that 802.1D-1990 specifies that the range for this parameter is related to the value of rcStgBridgeMaxAge. The granularity of this timer is specified by 802.1D-1990 to be 1 second. An agent can return a badValue error if a set is attempted to a value which is not a whole number of seconds.
DesignatedRoot	The unique identifier of the Bridge recorded as the Root in the Configuration BPDUs that are transmitted by the Designated Bridge for the segment to which the port is attached. Reference IEEE 802.1D-1990: Section 4.5.5.4.
RootCost	The cost of the path to the root as seen from this bridge
RootPort	The port number of the port which offers the lowest cost path from this bridge to the root bridge.
MaxAge	The maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded. The maximum age is specified in units of hundredths of a second. This is the actual value that bridge uses.
HelloTime	The amount of time required for transmission of the configuration BPDUs by the node on any port when it is the root of the spanning tree or trying to become the root. This is specified in units of hundredths of a second. This is the actual value that bridge uses.
ForwardDelay	This time value, measured in units of hundredths of a second, controls how fast a port changes its spanning state when moving towards the Forwarding state. The value determines how long the port stays in each of the Listening and Learning states, which precede the Forwarding state. This value is also used when a topology change has been detected, and is underway to age all dynamic entries in the Forwarding Database.
RstpUpCount	The number of times RSTP Module has been enabled. A Trap is generated on the occurrence of this event.
RstpDownCount	The number of times RSTP Module has been disabled. A Trap is generated on the occurrence of this event.
NewRootIdCount	The number of times this Bridge has detected a Root Identifier change. A Trap is generated on the occurrence of this event.

Variable	Value
TimeSinceTopologyChange	The time (in hundredths of a second) since the TcWhile Timer for any port in this Bridge was non-zero for Common Spanning Tree context.
TopChanges	The total number of topology changes detected by this bridge since the management entity was last reset or initialized.

RSTP Ports tab

Use the following procedure to view the **RSTP Ports** tab.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **RSTP**.
4. Select the **RSTP Ports** tab.

Variable definitions

Variable	Value
Port	The port number.
State	Every 2 bitfields used to identify a port state in this STG. Port state is cataloged as non-stp(0), blocking(1), learning(2), and forwarding(3).
Priority	The value of the priority field which is contained in the first (in network byte order) octet of the (2 octet long) Port ID.
PathCost	The contribution of this port to the path cost of paths towards the spanning tree root which include this port.
ProtocolMigration	The number of times this port has migrated from one STP protocol version to another. The relevant protocols are: • STP-COMPATIBLE • RSTP A Trap is generated on the occurrence of this event.
AdminEdgePort	The administrative value of the Edge Port parameter. A value of TRUE(1) indicates that this port should be assumed as an edge-port

Variable	Value
	and a value of FALSE(2) indicates that this port should be assumed as a non-edge-port.
OperEdgePort	The operational value of the Edge Port parameter. The object is initialized to FALSE on reception of a BPDU.
AdminPointToPoint	The administrative point-to-point status of the LAN segment attached to this port. A value of forceTrue(0) indicates that this port should always be treated as if it is connected to a point-to-point link. • A value of forceFalse or 1 indicates that this port should be treated as having a shared media connection. • A value of auto or 2 indicates that this port is considered to have a point-to-point link if it is an Aggregator and all of its members are aggregatable, or if the MAC entity is configured for full duplex operation, either through auto-negotiation or by management means.
OperPointToPoint	The operational point-to-point status of the LAN segment attached to this port. It indicates whether a port is considered to have a point-to-point connection or not. The value is determined by management or by auto-detection.
Participating	This field specifies whether a port is participating in the 802.1w protocol.
DesignatedRoot	The bridge identifier of the old root of the Spanning Tree as determined by the Spanning Tree Protocol as executed by this node.
DesignatedCost	The path cost of the Designated Port of the segment connected to this port. This value is compared to the Root Path Cost field in received BPDUs.
DesignatedBridge	The Bridge Identifier of the bridge which this port considers to be the Designated Bridge for this port's segment.
DesignatedPort	The Port Identifier for the port segment which is on the Designated Bridge.
ForwardTransitions	The number of times this port has transitioned from the Learning state to the Forwarding state.

RSTP Status tab

Use the following procedure to view the **RSTP Status** tab.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **RSTP**.
4. Select the **RSTP Status** tab.

Variable definitions

Variable	Value
Port	The port number.
Role	A role represents a functionality characteristic or capability of a resource to which policies are applied.
OperVersion	This indicates whether the Port is operationally in the RSTP mode or the STP-compatible mode for example, whether the Port is transmitting RST BPDUs or Config/TCN BPDUs.
EffectivePortState	This is the effective Operational state of the port. This object will be set to TRUE only when the port is operationally up in the interface manager and the force Port State for this port and specified port state is enabled. Otherwise this object is set to FALSE

Graphing RSTP Port Statistics using EDM

The **RSTP Status** tab shows RSTP Port statistics.

Use the following procedure to open the **RSTP Status** tab for graphing.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **RSTP**.
4. Select the **RSTP Status** tab.
5. Select a port and click on **Graph** to get the statistics for the selected port.

Variable definitions

Variable	Value
RxRstBpduCount	The number of RST BPDUs that were received on this port.
RxConfigBpduCount	The number of Configuration BPDUs that were received on this port.
RxTcnBpduCount	The number of TCN BPDUs that were received on this port.
TxRstBpduCount	Displays the number of RST BPDUs transmitted from this port.
TxConfigBpduCount	The number of Configuration BPDUs transmitted from this port.
TxTcnBpduCount	Signifies the number of TCN BPDUs transmitted from this port.
InvalidRstBpduRxCount	Number of Invalid RST BPDUs received on this port.
InvalidConfigBpduRxCount	Number of Invalid Configuration BPDUs received on this port.
InvalidTcnBpduRxCount	Number of Invalid TCN BPDUs received on this port.
ProtocolMigrationCount	The number of times this port has migrated from one STP protocol version to another. The relevant migration protocols are STP-COMPATIBLE and RSTP/MSTP. A trap is generated when the port migrates.

Multiple Spanning Tree Protocol

With Multiple Spanning Tree Protocol (MSTP or IEEE 802.1s), the user can configure multiple instances of RSTP on the same switch. Each RSTP instance can include one or more VLANs. The operation of the MSTP is similar to the current Avaya proprietary STG.

In the MSTP mode, the Ethernet Routing Switch 2500 Series supports a maximum of one Common and Internal Spanning Tree (CIST) and seven Multiple Spanning Tree Instances (MSTI).

Important:

You can access the MSTP menu command only when the switch is operating in the MSTP mode.

MSTP Globals tab

Use the following procedure to view the **MSTP Globals** tab.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **MSTP**.

   The MSTP dialog box with the **Globals** tab is displayed.

Variable definitions

Variable	Value
PathCostDefaultType	The version of the Spanning Tree default Path Costs that are to be used by this Bridge. A 16-bit value uses the 16-bit default path costs from IEEE Standard 802.1D-1998. A 32-bit value uses the 32-bit default path costs from IEEE Standard 802.1t.
TxHoldCount	The value used by the Port Transmit state machine to limit the maximum transmission rate.
MaxHopCount	The Maximum Hop Count value. The granularity of this timer is specified to be 1 second. An agent can return a badValue error if a set is attempted to a value which is not a whole number of seconds.
NoOfInstancesSupported	Indicates maximum number of spanning tree instances supported.
MstpUpCount	The number of times the MSTP Module has been enabled. A Trap is generated on the occurrence of this event.
MstpDownCount	The number of times the MSTP Module has been disabled. A Trap is generated on the occurrence of this event.
ForceProtocolVersion	Signifies the version of the Spanning Tree Protocol that the bridge is currently running.

Variable	Value
	• stpCompatible indicates that the bridge is using the Spanning Tree Protocol as specified in IEEE 802.1D. • rstp indicates that the bridge is using the Rapid Spanning Tree Protocol as specified in IEEE 802.1w • mstp indicates that the bridge is running the Multiple Spanning Tree Protocol as specified in IEEE 802.1s.
BrgAddress	The bridge address is generated when events like protocol up or protocol down occurs.
Root	The bridge identifier of the Root of the common spanning tree as determined by the Spanning Tree Protocol as executed by this node. This value is used as the CIST Root Identifier parameter in all Configuration BPDUs originated by this node.
RegionalRoot	The bridge identifier of the root of the Multiple spanning tree region as determined by the Spanning Tree Protocol as executed by this node. This value is used as the CIST Regional Root Identifier parameter in all Configuration Bridge PDUs originated by this node.
RootCost	The cost of the path to the CIST Root as seen from this bridge.
RegionalRootCost	The cost of the path to the CIST Regional Root as seen from this bridge.
RootPort	The port number of the port which offers the lowest path cost from the bridge to the CIST Root Bridge
BridgePriority	The value of the writable portion of the Bridge Identifier comprising of the first two octets. The values that are set for Bridge Priority must be in steps of 4096.
BridgeMaxAge	The value that all bridges use for MaxAge when this bridge is acting as the root. The granularity of this timer is specified to be 1 second. An agent can return a badValue error if a set is attempted to a value which is not a whole number of seconds.
BridgeForwardDelay	The value that all bridges use for ForwardDelay when this bridge is acting as the root. IEEE 802.1D specifies that the range for this parameter is related to the value of BridgeMaxAge. The granularity of this timer is specified to be 1 second. An agent can return a badValue error if a set is attempted to a value which is not a whole number of seconds.
HoldTime	This value determines the time interval during which no more than two Configuration BPDUs shall be transmitted

Variable	Value
	by this node. This value is measured in units of hundredths of a second
MaxAge	The maximum age of the Spanning Tree Protocol information learned from the network on any port before it is discarded. This value is measured in units of hundredths of a second.
ForwardDelay	This value controls how fast a port changes its spanning state when moving towards the Forwarding state. This value determines how long the port stays in a particular state before moving to the next state. It is measured in units of hundredths of a second.
TimeSinceTopology Change	The value (measured in hundredths of a second) The time since the TcWhile Timer for any port in this Bridge was non-zero for Common Spanning Tree context.
TopChanges	The number of times that there have been atleast one non-zero TcWhile Timer on this Bridge for the Common Spanning Tree context.
NewRootBridgeCount	The number of times this Bridge has detected a Root Bridge change for the Common Spanning Tree context. A Trap is generated when this event occurs.
RegionName	Signifies the name of the Region's configuration. By default, the Region Name is equal to the Bridge Mac Address.
RegionVersion	Denotes the version of the MST Region.
ConfigIdSel	The Configuration Identifier Format Selector used by the Bridge. This has a fixed value of 0 which is used to indicate RegionName, RegionVersion as specified in standard.
ConfigDigest	Signifies the Configuration Digest value for this Region. This is an MD5 digest value, and hence must always be 16 octets long.
RegionConfigChange Count	The number of times a Region Configuration Identifier Change was detected. A Trap is generated when this event occurs.

CIST Port tab

Use the following procedure to view the **CIST Port** tab.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **MSTP**.
4. Select the **CIST Port** tab.

Variable definitions

VariableValue	Value
Port	The port number of the port containing Spanning Tree information.
PathCost	The contribution of this port to the path cost of paths towards the CIST Root which include this port.
Priority	The four most significant bits of the Port Identifier of the Spanning Tree instance. It can be modified by setting the CistPortPriority value. The values that are set for Port Priority must be in steps of 16.
DesignatedRoot	This field specifies the unique Bridge Identifier of the bridge. It is recorded as the CIST Root in the configuration BPDUs which are transmitted.
DesignatedCost	The path cost of the Designated Port of the segment connected to this port.
DesignatedBridge	The unique Bridge Identifier of the bridge which the port considers to be the Designated Bridge for the port's segment.
DesignatedPort	The Port identifier of the port on the Designated Bridge which is designated for the port's segment.
RegionalRoot	Displays the unique Bridge Identifier of the bridge. It is recorded as the CIST Regional Root Identifier in the configuration BPDUs which are transmitted.
RegionalPathCost	The contribution of this port to the cost of paths. This value denotes the path of costs for the path towards the CIST Regional Root which include this port.
ProtocolMigration	Generated when port protocol migration happens in the port.
AdminEdgeStatus	The administrative value of the Edge Port parameter. A value of TRUE indicates that this port to be assumed as an edge-

VariableValue	Value
	port and a value of FALSE indicates that this port to be assumed as a non-edge-port.
OperEdgeStatus	Signifies the operational value of the Edge Port parameter. It is initialized to the value of AdminEdgeStatus and is set to FALSE when the port receives a BPDU.
AdminP2P	The administrative point-to-point status of the LAN segment attached to this port. A value of 0 indicates that this port should always be treated as if it is connected to a point-to-point link. A value of 1 indicates that this port should be treated as having a shared media connection. A value of 2 indicates that this port is considered to have a point-to-point link if it is an Aggregator and all of its members are aggregatable, or if the MAC entity is configured for full duplex operation, either through auto-negotiation, or by management means.
OperP2P	This field indicates the operational point-to-point status of the LAN segment attached to the port. It also indicates whether a port is considered to have a point-to-point connection or not. The value is determined by management or by auto-detection, as described in the AdminP2P object
HelloTime	The amount of time between the transmission of Configuration BPDUs transmitted by this node on the port. It is measured in units of hundredths of a second.
OperVersion	This indicates whether the port is operationally in the MSTP mode, RSTP mode or the STP-compatible mode for example, whether the port is transmitting MST BPDUs, RST BPDUs, or Config/TCN BPDUs.
EffectivePortState	The effective operational state of the port for CIST. This will be set to TRUE only when the port is operationally up in the Interface level and Protocol level for CIST. This is will be set to FALSE for all other times.
State	The current state of the port as defined by the Common Spanning Tree Protocol.
ForcePortState	The current state of the port which can be changed to either Disabled or Enabled for the base Spanning Tree instance.
SelectedPortRole	Selected port role of the port for the Spanning Tree instance.
CurrentPortRole	Current port role of the port for the Spanning Tree instance.

Graphing CIST Port Statistics

The **CIST Port** tab shows **CIST Port** statistics.

Use the following procedure to open the **CIST Port** tab for graphing.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **MSTP**.
4. Select the **CIST Port** tab.
5. Select a port and click on **Graph** to get the statistics for the CIST Port.

Variable definitions

Variable	Value
ForwardTransitions	The number of times this port has transitioned to the Forwarding State.
RxMstBpduCount	The number of MST BPDUs that have been received on this port.
RxRstBpduCount	The number of RST BPDUs that were received on this port.
RxConfigBpduCount	The number of Configuration BPDUs that were received on this port.
RxTcnBpduCount	The number of TCN BPDUs that were received on this port.
TxMstBpduCount	Shows the number of MST BPDUs transmitted from this port.
TxRstBpduCount	Displays the number of RST BPDUs transmitted from this port.
TxConfigBpduCount	The number of Configuration BPDUs transmitted from this port.
TxTcnBpduCount	Signifies the number of TCN BPDUs transmitted from this port.
InvalidMstBpduRxCount	Number of Invalid MST BPDUs received on this port.
InvalidRstBpduRxCount	Number of Invalid RST BPDUs received on this port.
InvalidConfigBpduRxCount	Number of Invalid Configuration BPDUs received on this port.
InvalidTcnBpduRxCount	Number of Invalid TCN BPDUs received on this port.
ProtocolMigrationCount	The number of times this port has migrated from one STP protocol version to another. The relevant migration

Variable	Value
	protocols are STP-COMPATIBLE and RSTP/MSTP. A trap is generated when the port migrates.

MSTI Bridges tab

Use the following procedure to view the **MSTI Bridges** tab.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **MSTP**.
4. Select the **MSTI Bridges** tab.

Variable definitions

Variable	Value
Instance	Spanning Tree Instance to which the information belongs.
RegionalRoot	Indicates MSTI Regional Root Identifier value for the Instance. This value is used as the MSTI Regional Root Identifier parameter in all Configuration Bridge PDUs originated by this node.
Priority	The writable portion of the MSTI Bridge Identifier comprising of the first two octets. The values that are set for Bridge Priority must be in steps of 4096.
RootCost	The cost of the path to the MSTI Regional Root as seen by this bridge.
RootPort	The port number of the port which offers the lowest path cost from this bridge to the MSTI Region Root Bridge.
Enabled	Used to control whether the bridge instance is enabled or disabled.
TimeSinceTopology Change	The time (measured in hundredths of a second) since the TcWhile Timer for any port in this bridge was non-zero for this Spanning Tree instance.

Variable	Value
TopChanges	The number of times that there have been at least one non-zero TcWhile Timer on this Bridge for this Spanning Tree instance.
NewRootCount	The number of times this bridge has detected a Root Bridge change for this Spanning Tree instance. A Trap is generated on the occurrence of this event.
InstanceUpCount	The number of times a new Spanning Tree instance has been created. A Trap is generated on the occurrence of this event.
InstanceDownCount	The number of times a Spanning Tree instance has been deleted. A Trap is generated on the occurrence of this event.

Inserting MSTI Bridges

Use the following procedure to insert MSTI Bridges.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **MSTP**.
4. Select the **MSTI Bridges** tab.
5. Click **Insert**.
6. Type the instance id.
7. Click **Insert**.

Deleting MSTI Bridges

Use the following procedure to delete MSTI Bridges.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **MSTP**.

4. Select the **MSTI Bridges** tab.
5. Click on one or multiple MSTI Bridges.
6. Click **Delete**.
7. To confirm you wish to delete the MSTI bridge, click **Yes**.

MSTI Port tab

Use the following procedure to view the **MSTI Port** tab.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. Double-click **Spanning Tree**.
3. Double-click **MSTP**.
4. Click the **MSTI Port** tab.

Variable definitions

Variable	Value
Port	Denotes the port number.
Instance	The number of times a Spanning Tree instance has been deleted. A Trap is generated when this event occurs.
State	Indicates the current state of the port as defined by the Multiple Spanning Tree Protocol. The state of a port can be Forwarding state in one instance, and Discarding (Blocking) state in another instance.
ForcePortState	Signifies the current state of the port which can be changed to either Disabled or Enabled for the specific Spanning Tree instance.
PathCost	The contribution of this port to the cost of paths towards the MSTI Root which includes this port.
Priority	Indicates the four most significant bits of the Port Identifier for a given Spanning Tree instance. It can be modified independently for each Spanning Tree instance supported by the bridge. The values that are set for Port Priority must be in steps of 16.

Variable	Value
DesignatedRoot	The unique Bridge Identifier of the bridge recorded as the MSTI Regional Root in the configuration BPDUs that are transmitted.
DesignatedBridge	The unique Bridge Identifier of the bridge which this port considers to be the Designated Bridge for the port's segment.
DesignatedPort	The Port identifier of the port on the Designated Bridge for this port's segment.
DesignatedCost	The path cost of the Designated Port of the segment connected to this port.
CurrentPortRole	Current Port Role of the port for this spanning tree instance.
EffectivePortState	The effective operational state of the port for specific instance. This is will be TRUE only when the port is operationally up in the interface level and Protocol level for the specific instance. This is will be set to FALSE at all other times.

Graphing MSTI port statistics using EDM

The **MSTI Port** tab can be used to graph MSTI port statistics.

Use the following procedure to open the **MSTI Port** tab for graphing.

Procedure steps

1. From the navigation tree, choose **VLAN , Spanning Tree** .
2. Double-click **MSTP**.
3. Select the **MSTI Port** tab.
4. Select a port and click on **Graph** to get the statistics for the MSTI Port.

Variable definitions

Variable	Value
ForwardTransitions	Number of times this port has transitioned to the Forwarding State for specific instance.
InvalidBPDUsRcvd	Number of Invalid BPDUs received on this Port for this Spanning Tree instance.

Variable	Value
ReceivedBPDUs	Number of BPDUs received by this port for this spanning tree instance.
TransmittedBPDUs	Number of BPDUs transmitted on this port for this Spanning Tree instance.

Setting up bridging

The Bridge parameters allow you to configure the global Spanning Tree and to view MAC address table for an Ethernet Routing Switch 2500 Series . Bridge information also includes Spanning Tree Group (STG) information.

Base tab

The MAC address used by the bridge must be referred to in a unique fashion; moreover, it should be the smallest MAC address (numerically) of all ports that belong to the bridge. However, it is only required to be unique when integrated with dot1dStpPriority. A unique BridgeIdentifier is formed that is used in the Spanning Tree Protocol.

Use the following procedure to view the **Base** tab.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. Double-click **Bridge**.
3. In the work area, click the **Base** tab.

Variable definitions

Variable	Value
BridgeAddress	MAC address of the bridge when it is referred to in a unique fashion. This address should be the smallest MAC address of all ports that belong to the bridge. However, it is has to be unique. When concatenated with dot1dStpPriority, a unique bridge ID is formed that is then used in the Spanning Tree Protocol.
NumPorts	Number of ports controlled by the bridging entity.

Variable	Value
Type	Indicates the type of bridging this bridge can perform. If the bridge is actually performing a certain type of bridging, this will be indicated by entries in the port table for the given type.

Transparent tab

The **Transparent** tab contains information about a specific unicast MAC address that has forwarding information for the bridge.

Use the following procedure to view the **Transparent** tab.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. Double-click **Bridge**.
3. Select the **Transparent** tab.

Variable definitions

Use the data in the following table to help you understand the transparent tab.

Variable	Value
LearnedEntryDiscards	Number of Forwarding database entries learned that have been discarded due to a lack of space in the Forwarding database. If this counter is increasing, it indicates that the Forwarding database is becoming full regularly. This condition will effect the performance of the subnetwork. If the counter has a significant value and is not presently increasing, it indicates that the problem has been occurring but is not persistent.
AgingTime	Time-out period in seconds for aging out dynamically learned forwarding information.  **Important:** The 802.1D-1990 specification recommends a default of 300 seconds.

Forwarding tab

The **Forwarding** tab displays the current MAC Address Table (Forwarding table) on the switch.

Use the following procedure to view the **Forwarding** tab.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. Double-click **Bridge**.
3. Select the **Forwarding** tab.

Variable definitions

Use the data in the following table to help you understand the MAC Address Table (Forwarding table).

Variable	Value
Id	Specifies the VLAN identifier.
Address	A unicast MAC address for which the bridge has forwarding or filtering information.
Port	Either the value "0" or the port number on a frame has been seen. The source address must be equal to the value of the corresponding instance of dot1dTpFdbAddress. A value of "0" indicates that the port number has not been learned, so the bridge does have the forwarding/filtering information for this address (located in the dot1dStaticTable). You should assign the port value to this object whenever it is learned even for addresses for which the corresponding value of dot1dTpFdbStatus is not learned(3).
Status	The values of this fields include: • invalid: Entry is no longer valid, but has not been removed from the table. • learned: Value of the corresponding instance of dot1dTpFdbPort was learned and is being used. • self: Value of the corresponding instance of dot1dTpFdbAddress represents an address of the bridge. The corresponding instance of dot1dTpFdbPort indicates that a specific port on the bridge has this address.

Variable	Value
	• mgmt(5): Value of the corresponding instance of dot1dTpFdbAddress is also the value of an existing instance of dot1dStaticAddress. • other: none of the preceding. This would include where some other MIB object (not the corresponding instance of dot1dTpFdbPort or an entry in the dot1dStaticTable) is being used to determine if a frames addressed to the value of dot1dTpFdbAddress are being forwarded.

Graphing port bridge statistics

Use the following procedure to graph port bridge statistical information.

Procedure steps

1. From the Device Physical View, click a port.
2. From the navigation tree, double-click **Graph**.
3. In the Graph tree, double-click **Port** .
4. In the work area, click the **Bridge** tab.
5. Click the down arrow to the left of the **Poll Interval** dialog box.
6. Select a value from the list.
7. To reset the statistics counters, click **Clear Counters**.
8. To select bridge statistical information to graph, click an information row.
9. Click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart** column.

Variable definitions

Use the data in the following table to graph port bridge statistics.

Variable	Value
DelayExceededDiscards	Number of frames discarded by the port due to excessive transit delays through the bridge. It is incremented by both transparent and source route bridges.
MtuExceededDiscards	Number of frames discarded by the port due to an excessive size. It is incremented by both transparent and source route bridges.

Variable	Value
InFrames	The number of frames that have been received by this port from its segment.
OutFrames	The number of frames that have been received by this port from its segment.
InDiscards	Count of valid frames received which were discarded (filtered) by the Forwarding Process.

Chapter 14: Configuring Multi-Link Trunking using Enterprise Device Manager

Multi-Link Trunking (MLT) is a point-to-point connection that aggregates multiple ports so that they logically act like a single port with the aggregated bandwidth. You can achieve higher aggregate throughput on a switch-to-switch or switch-to-server application by grouping multiple ports into a logical link . Multi-Link Trunking provides media and module redundancy.

The chapter includes the following information:

Multi-Link Trunk features

A number of Avaya products implement Multi-Link Trunking (MLT) and have different features and requirements based on the architecture of the device. For the Ethernet Routing Switch 2500 Series , Multi-Link Trunking has the following general features and requirements:

- A unit can have up to six Multi-Link Trunks (MLTs).
- Up to four ports can belong to an MLT.
- Multi-Link Trunking is supported on 10BASE-T, 100BASE-TX, and SFP ports.
- Multi-Link Trunking is compatible with the Spanning Tree Protocol.
- IEEE 802.1Q tagging is supported on an MLT.
- The distribution algorithm is user-programmable. The default algorithm that distributes traffic across an MLT is based on the source and destination MAC addresses (BASIC mode). An algorithm that distributes traffic based on the source and destination IP addresses (ADVANCE mode) is also available.
- Distributed MLT (D-MLT) support. D-MLT is MLT with ports from two or more stack units.

Configuring Multi-Link Trunks using EDM

Use the following procedure to display and configure MLTs using EDM.

Procedure steps

1. From the navigation tree, choose **VLAN**.
2. Double-click **MLT/LACP**.
3. In the work area, click the **Multi-Link Trunks** tab.
4. To select a trunk to create, click the trunk **Id**.
5. In the trunk row, double-click the cell in the **Name** column.
6. In the field, type a name for the MLT, or accept the default name.
7. In the trunk row, double-click the cell in the **PortMembers** column.
8. From the list, select multiple ports to add to the trunk.
9. Click **Ok**.
10. In the trunk row, double-click the cell in the **Loadbalance(Mode)** column.
11. From the list, select a load balancing mode.
12. In the trunk row, double-click in the **Enable** column.
13. From the list, select **true** to enable the MLT, or **false** to disable the MLT.
14. To create additional MLTs, repeat steps 4 to 13.
15. Click **Apply**.

Variable definitions

Variable	Value
Id	The MLT identification number (assigned consecutively).
PortType	Access or trunk port.
Name	The name given to the MLT.
PortMembers	The ports assigned to the MLT.
VlanIds	The VLANs assigned to the MLT.
Loadbalance(Mode)	Specifies the load balance mode. Values include: • basic • advanced
Enable	Specifies enabling of the MLT.

Viewing MLT utilization using EDM

Use the following procedure to display MLT utilization information.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **MLT/LACP**.
3. In the work area, click the **MLT Utilization** tab.

Variable definitions

Variable	Value
MltId	Specifies the MLT Identification number.
PortIfIndex	Specifies the port identification number.
TrafficType	Specifies the traffic type.
TrafficLast5Min	Specifies the MLT traffic in the last five minutes.
TrafficLast30Min	Specifies the MLT traffic in the last thirty minutes.
TrafficLast1Hour	Specifies the MLT traffic in the last hour.

Graphing Multi-Link Trunk statistics using EDM

Use the following procedure to display and graph MLT interface statistics.

Procedure steps

1. From the navigation tree, choose **VLAN**.
2. Double-click **MLT/LACP**.
3. In the work area, click the **MultiLink Trunks** tab.
4. To select an MLT to graph, click the trunk Id.
5. Click **Graph**.
6. Click the **Interface** tab.
7. Select a **Poll Interval** from the list.
8. From the list, select a poll interval time.
9. To reset the MLT statistic counters, click **Clear Counters**.

10. To select statistics to graph, click a statistic type row under one of the displayed columns.
11. Click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.
12. To return to the MutiLink Trunks—Graph work area, click **Close**.

Variable definitions

Variable	Value
InMulticastPkts	The number of packets delivered to this MLT that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both Group and Functional addresses.
OutMulticastPkts	The total number of packets that higher-level protocols requested to be transmitted, and that were addressed to a multicast address at this MLT, including those that were discarded or not sent. For a MAC layer protocol, this number includes both Group and Functional addresses.
InBroadcastPkts	The number of packets delivered to this MLT that were addressed to a broadcast address at this sublayer.
OutBroadcastPkts	The total number of packets that higher-level protocols requested to be transmitted, and that were addressed to a broadcast address at this MLT, including those that were discarded or not sent.
HCInOctets	The total number of octets received on the MLT interface, including framing characters.
HCOutOctets	The total number of octets transmitted out of the MLT interface, including framing characters.
HCInUcastPkts	The number of packets delivered by this MLT to a higher MLT that were not addressed to a multicast or broadcast address at this sublayer.
HCOutUcastPkts	The number of packets that high-level protocols requested to be transmitted that were not addressed to a multicast address at this MLT. This total number includes those packets discarded or unsent.
HCInMulticastPkt	The number of packets delivered to this MLT that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both Group and Functional addresses.
HCOutMulticast	The total number of packets that high-level protocols requested to be transmitted, and that were addressed to a multicast address at this MLT, including those that were discarded or not

Variable	Value
	sent. For a MAC layer protocol, this number includes both Group and Functional addresses.
HCInBroadcastPkt	The number of packets delivered to this MLT that were addressed to a broadcast address at this sublayer.
HCOutBroadcast	The total number of packets that high-level protocols requested to be transmitted, and that were addressed to a broadcast address at this MLT, including those that were discarded or not sent.

Graphing Multi-Link Trunk Ethernet error statistics using EDM

Use the following procedure to display and graph Multi-Link Trunk Ethernet error statistics.

Procedure steps

1. From the navigation tree, choose **VLAN**.
2. Double-click **MLT/LACP**.
3. In the work area, click the **MultiLink Trunks** tab.
4. To select an MLT to graph, click the trunk Id.
5. Click **Graph**.
6. Click the **Ethernet Errors** tab.
7. Select a **Poll Interval** from the list.
8. From the list, select a poll interval time.
9. To reset the MLT statistic counters, click **Clear Counters**.
10. To select statistics to graph, click a statistic type row under one of the displayed columns.
11. Click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.
12. To return to the MultiLink Trunks-Graph tab, click **Close**.

Variable definitions

Variable	Value
AlignmentErrors	A count of frames received on a particular MLT that are not an integral number of octets in length and do not pass the FCS check.

Variable	Value
	The count represented by an instance of this object is incremented when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
FCSErrors	A count of frames received on an MLT that are an integral number of octets in length but do not pass the FCS check. The count represented by an instance of this object is incremented when the frameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
IMacTransmit Error	A count of frames for which transmission on a particular MLT fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the LateCollisions object, the ExcessiveCollisions object, or the CarrierSenseErrors object.
IMacReceive Error	A count of frames for which reception on a particular MLT fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the FrameTooLongs object, the AlignmentErrors object, or the FCSErrors object. The precise meaning of the count represented by an instance of this object is implementation specific. In particular, an instance of this object can represent a count of received errors on a particular interface that are not otherwise counted.
CarrierSense Error	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular MLT. The count represented by an instance of this object is incremented at most once per transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
FrameTooLong	A count of frames received on a particular MLT that exceed the maximum permitted frame size. The count represented by an instance of this object is incremented when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
SQETestError	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular MLT. The SQE TEST ERROR message is defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation is described in section 7.2.4.6 of the same document.

Variable	Value
Deferred Transmiss	A count of frames for which the first transmission attempt on a particular MLT is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.
SingleCollFrames	A count of successfully transmitted frames on a particular MLT for which transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts, ifOutMulticastPkts, or ifOutBroadcastPkts, and is not counted by the corresponding instance of the MultipleCollisionFrames object.
MultipleColl Frames	A count of successfully transmitted frames on a particular MLT for which transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts, ifOutMulticastPkts, or ifOutBroadcastPkts, and is not counted by the corresponding instance of the SingleCollisionFrames object.
LateCollisions	The number of times that a collision is detected on a particular MLT later than 512 bit-times into the transmission of a packet. Five hundred and twelve bit-times corresponds to 51.2 microseconds on a 10 Mb/s system. A (late) collision included in a count represented by an instance of this object is also considered as a (generic) collision for purposes of other collision-related statistics.
ExcessiveCollis	A count of frames for which transmission on a particular MLT fails due to excessive collisions.

Link Aggregation Control Protocol

With Link Aggregation (LA), you can create and manage a trunk group. You can control and configure a trunk group automatically through the use of the Link Aggregation Control Protocol (LACP).

The LACP, defined by the IEEE 802.3ad standard, allows a switch to learn the presence and capabilities of a remote switch by exchanging information with the remote switch before a trunk group is formed. Either switch can accept or reject the aggregation request with the far end on a per port basis. A link that can not join a trunk group operates as an individual link.

Viewing LAG information using EDM

Use the following procedure to display Link Aggregation Group (LAG) configuration information.

Procedure steps

1. From the navigation tree, choose **VLAN**.
2. Double-click **MLT/LACP**.
3. Select the **LACP** tab.

Variable definitions

Variable	Value
Index	The unique identifier allocated to this Aggregator by the local System. This attribute identifies an Aggregator instance among the subordinate managed objects of the containing object. This value is read-only.
MacAddress	The MAC address used by this bridge when it must be referred to in a unique fashion.
AggregateOrIndividual	A read-only Boolean value indicating whether the Aggregation Port is able to Aggregate ('TRUE') or is only able to operate as an Individual link ('FALSE').
ActorLagId	The combined information of ActorSystemPriority, ActorSystemID, and ActorOperKey in "ActorSystemPriority-ActorSystemID-ActorOperKey" format.
ActorSystemPriority	A 2-octet read-write value used to define the priority value associated with the Actor's System ID.
ActorSystemID	A 6-octet read-only MAC address value that defines the value of the System ID for the System that contains this Aggregation Port.
ActorOperKey	The current operational value of the Key for the Aggregation Port. This is a 16-bit read-only value.
ActorAdminKey	The current administrative value of the Key for the Aggregation Port. This is a 16-bit read-write value.
PartnerLagId	The combined information of PartnerSystemPriority, PartnerSystemID, and PartnerOperKey in "PartnerSystemPriority-PartnerSystemID-PartnerOperKey" format.
PartnerSystemPriority	A 2-octet read-only value that indicates the priority value associated with the Partner's System ID.

Variable	Value
PartnerSystemID	A 6-octet read-only MAC address value consisting of the unique identifier for the current protocol Partner of this Aggregator. A value of zero indicates that there is no known Partner. If the aggregation is manually configured, this System ID value will be a value assigned by the local System.
PartnerOperKey	The current operational value of the Key for the Aggregator's current protocol Partner. This is a 16-bit read-only value.
CollectorMaxDelay	The value of this 16-bit read-write attribute defines the maximum delay, in tens of microseconds, that can be imposed by the Frame Collector between receiving a frame from an Aggregator Parser, and either delivering the frame to its MAC Client or discarding the frame

Link Aggregation Group configuration using EDM

Use the procedures in this section to display or modify LAG member configuration.

Viewing LACP for LAG members using EDM

Use the following procedure to display the existing LACP configuration for LAG members.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **MLT/LACP**.
3. In the work area, click the **LACP Ports** tab.

Variable Definitions

Use the data in this table to help you understand the LACP configuration for LAG members display.

Variable	Value
Index	Indicates the unique identifier allocated to an Aggregator by the local system.
AdminEnabled	Indicates the current administrative setting for the port. A value of true enables the port to participate in LACP. A value of false disables the port from participating in LACP.

Variable	Value
OperEnabled	Indicates the current operational state for the port. A value of true means the port is participating in LACP. A value of false means the port is not participating in LACP.
ActorAdminState	Indicates the Actor administrative state for the port. Values include: • lacpActive • aggregation • shortTimeout
ActorOperState	Indicates the current operational values of Actor state transmitted by the Actor in LACPDUs.
AggregateOrIndividual	Indicates whether the port represents an Aggregate or an Individual link.
ActorPortPriority	Indicates the priority value assigned to this Aggregation port. Values range from 0–65535.
ActorAdminKey	Indicates the current administrative value of the Key for the Aggregation Port. Values range from 1–4095.
ActorOperKey	Indicates the current operational value of the Key for the Aggregation Port.
SelectedAggID	Indicates the identifier value of the Aggregator that this Aggregation Port has currently selected. Zero indicates that the Aggregation Port has not selected an Aggregator, either because it is in the process of detaching from an Aggregator or because no suitable Aggregator exists for it to select.
AttachedAggID	Indicates the identifier value of the Aggregator that this Aggregation Port is currently attached to. Zero indicates that the Aggregation Port is not currently attached to an Aggregator. This value is read-only.
ActorPort	Indicates the port number locally assigned to the Aggregation Port. The port number is communicated in LACPDUs as the Actor_Port. This value is read-only
MltId	Indicates the MLT that the port is assigned to. If the port is not assigned to an MLT, the MltId value is 0.
PartnerOperPort	Indicates the operational port number assigned by the port protocol partner.
OperStatus	Indicates the operational status of the interface. Values are up (operational) or down (not operational).

Configuring LACP for specific LAG members using EDM

Use the following procedure to configure LACP for LAG members.

Prerequisites

- Ensure members you want to configure are not ADAC Call Server or Uplink ports.
- Disable ADAC for members you want to configure

Important:

To configure the port LACP mode to active, you must set the AdminEnabled value to **true** and the ActorAdminState value to **lacpActive.**

Important:

To configure the port LACP mode to passive, you must set the AdminEnabled value to **false** and clear the **lacpActive**, **aggregation**, and **shortTimeout** check boxes in ActorAdminState.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **MLT/LACP**.
3. In the work area, click the **LACP Ports** tab.
4. To select a port to configure, click the port **Index**.
5. In the port row, double-click the cell in the **AdminEnabled** column.
6. Select a value from the list—**true** to enable LACP for the port, or **false** to disable LACP for the port.
7. In the port row, double-click the cell in the **ActorAdminState** column.
8. Select an individual or combination of check boxes.
9. Click **Ok**.
10. In the port row, double-click the cell in the **ActorPortPriority** column.
11. In the dialog box, edit the value as required.
12. In the port row, double-click the cell in the **ActorAdminKey** column.
13. In the dialog box, edit the value as required.
14. Click **Apply**.

Variable Definitions

Use the data in this table to configure LACP for LAG members.

Variable	Value
Index	Indicates the unique identifier allocated to an Aggregator by the local system. This is a read-only cell.
AdminEnabled	Specifies the current administrative setting for the port. A value of true enables the port to participate in LACP. A value of false disables the port from participating in LACP. **Important:** You cannot enable ports to participate in LACP if they are members of an enabled MLT.
OperEnabled	Indicates the current operational state for the port. A value of true means the port is participating in LACP. A value of false means the port is not participating in LACP. This is a read-only cell.
ActorAdminState	Specifies the Actor administrative state. Values include: • lacpActive • aggregation • shortTimeout
ActorOperState	Indicates the current Actor operational state. This is a read-only cell.
AggregateOrIndividual	Indicates whether the Aggregator represents an Aggregate or an Individual link. This is a read-only cell.
ActorPortPriority	Specifies the priority value assigned to this Aggregation port. Values range from 0–65535.
ActorAdminKey	Specifies the current administrative value of the Key for the Aggregation Port. Values range from 1–4095.
ActorOperKey	Indicates the current operational value of the Key for the Aggregation Port. This is a read-only cell.
SelectedAggID	Indicates the identifier value of the Aggregator that this Aggregation Port has currently selected. Zero indicates that the Aggregation Port has not selected an Aggregator, either because it is in the process of detaching from an Aggregator or because no suitable Aggregator exists for it to select. This is a read-only cell.
AttachedAggID	Indicates the identifier value of the Aggregator that this Aggregation Port is currently attached to. Zero indicates that the Aggregation Port is not currently attached to an Aggregator. This is a read-only cell.
ActorPort	Indicates the port number locally assigned to the Aggregation Port. The port number is communicated in LACPDUs as the Actor_Port. This is a read-only cell.

Variable	Value
MltId	Indicates the MLT that the port is assigned to. If the port is not assigned to an MLT, the MltId value is 0. This is a read-only cell.
PartnerOperPort	The operational port number assigned by the port's protocol partner. This is a read-only cell.
OperStatus	Indicates the operational status of the interface. Values are up (operational) or down (not operational). This is a read-only cell.

LACP configuration for ports using EDM

You can use the information in this section to display or modify the LACP configuration for switch ports.

Viewing the LACP configuration for ports using EDM

Use the following procedure to display the existing LACP configuration for switch ports.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Ports**.
4. Click the **LACP** tab.

Variable definitions

Use the data in this table to help you understand the LACP configuration display for switch ports.

Variable	Value
Index	Indicates the unique identifier allocated to an Aggregator by the local system. This is a read-only cell.
AdminEnabled	Specifies the current administrative setting for the port. A value of true enables the port to participate in LACP. A value of false disables the port from participating in LACP. **Important:** You cannot enable ports to participate in LACP if they are members of an enabled MLT.

Variable	Value
OperEnabled	Indicates the current operational state for the port. A value of true means the port is participating in LACP. A value of false means the port is not participating in LACP. This is a read-only cell.
ActorAdminState	Specifies the Actor administrative state. Values include: • lacpActive • aggregation • shortTimeout
ActorOperState	Indicates the current Actor operational state. This is a read-only cell.
AggregateOrIndividual	Indicates whether the Aggregator represents an Aggregate or an Individual link. This is a read-only cell.
ActorPortPriority	Specifies the priority value assigned to this Aggregation port. Values range from 0–65535.
ActorAdminKey	Specifies the current administrative value of the Key for the Aggregation Port. Values range from 1–4095.
ActorOperKey	Indicates the current operational value of the Key for the Aggregation Port. This is a read-only cell.
SelectedAggID	Indicates the identifier value of the Aggregator that this Aggregation Port has currently selected. Zero indicates that the Aggregation Port has not selected an Aggregator, either because it is in the process of detaching from an Aggregator or because no suitable Aggregator exists for it to select. This is a read-only cell.
AttachedAggID	Indicates the identifier value of the Aggregator that this Aggregation Port is currently attached to. Zero indicates that the Aggregation Port is not currently attached to an Aggregator. This is a read-only cell.
ActorPort	Indicates the port number locally assigned to the Aggregation Port. The port number is communicated in LACPDUs as the Actor_Port. This is a read-only cell.
MltId	Indicates the MLT that the port is assigned to. If the port is not assigned to an MLT, the MltId value is 0.
PartnerOperPort	Indicates the operational port number assigned by the port's protocol partner. This is a read-only cell.
OperStatus	Indicates the operational status of the interface. Values are up (operational) or down (not operational).

Configuring LACP for specific ports using EDM

Use the following procedure to modify the LACP configuration for one or more switch ports.

Prerequisites

- Ensure ports you want to configure are not ADAC Call Server or Uplink ports.
- Disable ADAC for ports you want to configure

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Ports**.
4. Click the LACP tab.
5. To select a port to configure, click the port **Index**.
6. In the port row, double-click the cell in the **AdminEnabled** column.
7. Select a value from the list—**true** to enable LACP for the port, or **false** to disable LACP for the port.
8. In the port row, double-click the cell in the **ActorAdminState** column.
9. Select an individual or combination of check boxes
10. Click **Ok**.
11. In the port row, double-click the cell in the **ActorPortPriority** column.
12. In the dialog box, edit the value as required.
13. In the port row, double-click the cell in the **ActorAdminKey** column.
14. In the dialog box, edit the value as required.
15. Repeat steps 5 through 14 to configure LACP for additional ports as required.
16. Click **Apply**.

Variable definitions

Use the data in this table to configure LACP for specific ports.

Variable	Value
Index	Indicates the unique identifier allocated to an Aggregator by the local system. This is a read-only cell.
ActorSystemPriority	Specifies the priority value associated with the Actor System ID. Values range from 0–65535.
AdminEnabled	Specifies the current administrative setting for the port. A value of true enables the port to participate in LACP. A

Variable	Value
	value of false disables the port from participating in LACP. **Important:** You cannot enable ports to participate in LACP if they are members of an enabled MLT.
OperEnabled	Indicates the current operational state for the port. A value of true means the port is participating in LACP. A value of false means the port is not participating in LACP. This is a read-only cell.
ActorAdminState	Specifies the Actor administrative state. Values include: • lacpActive • aggregation • shortTimeout
ActorOperState	Indicates the current Actor operational state. This is a read-only cell.
AggregateOrIndividual	Indicates whether the Aggregator represents an Aggregate or an Individual link. This is a read-only cell.
ActorPortPriority	Specifies the priority value assigned to this Aggregation port. Values range from 0–65535.
ActorAdminKey	Specifies the current administrative value of the Key for the Aggregation Port. Values range from 1–4095.
ActorOperKey	Indicates the current operational value of the Key for the Aggregation Port. This is a read-only cell.
SelectedAggID	Indicates the identifier value of the Aggregator that this Aggregation Port has currently selected. Zero indicates that the Aggregation Port has not selected an Aggregator, either because it is in the process of detaching from an Aggregator or because no suitable Aggregator exists for it to select. This is a read-only cell.
AttachedAggID	Indicates the identifier value of the Aggregator that this Aggregation Port is currently attached to. Zero indicates that the Aggregation Port is not currently attached to an Aggregator. This is a read-only cell.
ActorPort	Indicates the port number locally assigned to the Aggregation Port. The port number is communicated in LACPDUs as the Actor_Port. This is a read-only cell.
MltId	Indicates the MLT that the port is assigned to. If the port is not assigned to an MLT, the MltId value is 0.

Variable	Value
PartnerOperPort	Indicates the operational port number assigned by the protocol partner of port. This is a read-only cell.
OperStatus	Indicates the operational status of the interface. Values are up (operational) or down (not operational).
Important: To configure the port LACP mode to active, you must set the AdminEnabled value to **true** and the ActorAdminState value to **lacpActive.**	
Important: To configure the port LACP mode to passive, you must set the AdminEnabled value to **false** and clear the **lacpActive**, **aggregation**, and **shortTimeout** check boxes in ActorAdminState.	

Graphing port LACP statistics using EDM

Use the following procedure to display and graph LACP statistics for switch ports.

Procedure steps

1. From the Device Physical View, click a port.
2. From the navigation tree, double-click **Graph**.
3. In the Graph tree, double-click **Port** .
4. In the work area, click the **LACP** tab.
5. Select a **Poll Interval** from the list.
6. To select statistics to graph, click a statistic type row under one of the displayed columns.
7. Click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Variable	Value
LACPDUsRx	Denotes the number of valid LACPDUs received on this Aggregation Port. This value is read-only.
MarkerPDUsRx	Signifies the number of valid Marker PDUs received on this Aggregation Port. This value is read-only.

Variable	Value
MarkerResponse PDUsRx	The number of valid Marker Response PDUs received on this Aggregation Port. This value is read-only.
UnknownRx	Indicates the number of frames received that can • Carry the Slow Protocols Ethernet Type value (43B.4), but contain an unknown PDU. • Are addressed to the Slow Protocols group MAC Address (43B.3), but do not carry the Slow Protocols Ethernet Type. This value is read-only.
IllegalRx	Denotes the number of frames received that carry the Slow Protocols Ethernet Type value (43B.4), but contain a badly formed PDU or an illegal value of Protocol Subtype (43B.4). This value is read-only.
LACPDUsTx	Signifies the number of LACPDUs that are transmitted on this Aggregation Port. This value is read-only.
MarkerPDUsTx	Displays the number of Marker PDUs transmitted on this Aggregation Port. This value is read-only.
MarkerResponse PDUsTx	Indicates the number of Marker Response PDUs that are transmitted on this Aggregation Port. This value is read-only.

Configuring MLT and VLACP global settings using EDM

Use the information in this section to

- enable or disable VLACP globally
- set the VLACP Multicast MAC Address
- enable or disable MLT whole trunk mode globally

Configuring MLT whole trunk using EDM

Use the following procedure to configure the MLT whole trunk mode of a switch or stack.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **MLT/LACP**.
3. On the work area, click the **Global** tab.

4. Select **MltDisablePortsOnShutdown** to enable or disable the MLT whole trunk feature.
5. On the toolbar, click **Apply**.

Enabling or disabling global VLACP using EDM

Use the following procedure to enable or disable VLACP for the switch.

Procedure steps

1. From the navigation tree, click **VLAN**.
2. In the VLAN tree, click **MLT/LACP**.
3. In the work area, click the **Global** tab.
4. Do one of the following:
   - To enable VLACP, select the **VlacpEnable** check box.
   - To disable VLACP, deselect the **VlacpEnable** check box.
5. Type a value in the **VlacpMulticastMACAddress** dialog box.
6. On the toolbar, click **Apply**.

Variable definitions

Field	Description
VlacpEnable	Enables or disables VLACP on the switch.
MulticastMACAddress	Identifies a multicast MAC address used exclusively for VLACPDUs. Default is 01:80:c2:00:11:00.

VLACP configuration for ports using EDM

Use the procedures in this section to view and configure VLACP at the port level.

Viewing the VLACP configuration for ports using EDM

Use the following procedure to view the VLACP tab for ports.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Ports**.
4. Click the **VLACP** tab.

Variable definitions

Variable	Value
rcPortIndex	Indicates the switch and port number.
AdminEnable	Enables or disables VLACP on a port. The default value is False.
OperEnable	Indicates whether VLACP is operationally enabled or disabled. This is a read-only field.
FastPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using short timeouts. Valid values range from 400-20000 with a default of 500.
SlowPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using long timeouts. Valid values range from 10000-30000 with a default of 30000.
Timeout	Specifies whether the timeout control value is a short or long timeout.
TimeoutScale	Sets a timeout scale for the port, where timeout = (periodic time) * (timeout scale). The range is 1-10. Default is 3. Note: With VLACP, a short interval exists between a port transmitting a VLACPDU and the partner port receiving the same VLACPDU. However, if the timeout-scale is set to 1, the port timeout value does not take into account the normal travel time of the VLACPDU. The port expects to receive a VLACPDU at the same moment the partner port sends it. Therefore, the delayed VLACPDU results in the link being blocked, and then enabled again when the packet arrives. To prevent this scenario from happening, set the timeout-scale to a value larger than 1.
EtherType	Specifies VLACP protocol identification. The ID value is a 4-digit Hex number, with a default of 8103.

Variable	Value
EtherMacAddress	The default value is 00:00:00:00:00:00 and it can be configured with the MAC address of the switch or stack to which this port is sending VLACPDUs. It cannot be configured as a multicast MAC. Note: VLACP has only one multicast MAC address, configured using the MulticastMACAddress field in the VLACP Global tab, which is the Layer 2 destination address used for the VLACPDUs. The port-specific EtherMACAddresss parameter does not specify a multicast MAC address, but instead specifies the MAC address of the switch or stack to which this port is sending VLACPDUs. You are not always required to configure EtherMACAddresss. If not configured, the first VLACP-enabled switch that receives the PDUs from a unit assumes that it is the intended recipient and processes the PDUs accordingly. If you want an intermediate switch to drop VLACP packets, configure the EtherMACAddresss field with the desired destination MAC address. With EtherMACAddresss configured, the intermediate switches do not misinterpret the VLACP packets.
PortState	Identifies whether the VLACP port state is up or down. This is a read-only field.

Configuring VLACP for specific ports using EDM

Use the following procedure to configure VLACP for a single port or multiple ports.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Ports**.
4. Click the **VLACP** tab.
5. To select a port to edit, click **rcPortIndex** row.
6. In the port row, double-click the cell in the **AdminEnable** column.
7. Select a value from the list—true to enable VLACP for the port, or false to disable VLACP for the port.
8. In the port row, double-click the cell in the **FastPeriodicTimer** column.
9. Type a value in the dialog box.

10. In the port row, double-click the cell in the **SlowPeriodicTimer** column.
11. Type a value in the dialog box.
12. In the port row, double-click the cell in the **Timeout** column.
13. Select a value from the list.
14. In the port row, double-click the cell in the **TimeoutScale** column.
15. Type a value in the dialog box.
16. In the port row, double-click the cell in the **EtherType** column.
17. Type a value in the dialog box.
18. In the port row, double-click the cell in the **EtherMacAddress** column.
19. Type a value in the dialog box.
20. Repeat steps **5** through **19** to configure VLACP for additional ports as required.
21. Click **Apply**.

Variable Definitions

Use the data in this table to edit the VLACP configuration for individual ports.

Variable	Value
rcPortIndex	Specifies the switch and port number.
AdminEnable	Indicates whether VLACP is enabled (true) or disabled (false) on ports. The default value is disabled.
OperEnable	Indicates whether VLACP is operationally enabled or disabled. This is a read-only cell. **Important:** VLACP is only operational when OperEnable is true and PortState is up.
FastPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using short timeouts. Valid values range from 400-20000 with a default of 500.
SlowPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using long timeouts. Valid values range from 10000-30000 with a default of 30000.
Timeout	Specifies whether the timeout control value is a short or long timeout.
TimeoutScale	Specifies the scale value used to calculate timeout from periodic time. Values range from 1–10. The default is 3. With VLACP, a short interval exists between a port transmitting a VLACPDU and the partner port receiving the

Variable	Value
	same VLACPDU. If the timeout-scale is set to 1, the port timeout value does not take into account the normal travel time of the VLACPDU. The port expects to receive a VLACPDU at the same moment the partner port sends it. Therefore, the delayed VLACPDU results in the link being blocked, and then enabled again when the packet arrives. Avaya recommends that you set the timeout scale to a value larger than 1.
EtherType	Specifies VLACP protocol identification. The value can be entered as a numerical value ranging from 33025–33279 or a hexadecimal equivalent (8101–81ff). The default is 8103. Use the prefix 0x to type a hexadecimal value in the dialog box. Only hexadecimal values display in the EtherType column of the VLACP work area.
EtherMacAddress	Specifies the MAC address of the switch or stack to which a port is sending VLACPDUs. The default value is 00:00:00:00:00:00. It cannot be configured as a multicast MAC. VLACP uses only the multicast MAC address configured when VLACP is enabled globally. This is the Layer 2 destination address used for the VLACPDUs. If you do not type a value for the EtherMACAddress, the first VLACP-enabled switch that receives the PDUs from a sending port becomes the intended recipient and processes the PDUs. If you want an intermediate switch to drop VLACP packets, configure EtherMACAddress with the desired destination MAC address. With EtherMACAddress configured, the intermediate switches do not misinterpret the VLACP packets.
PortState	Indicates whether the VLACP port state is up or down. This is a read-only cell. **Important:** VLACP is only operational when OperEnable is true and PortState is up.

Chapter 15: Configuring ADAC for Avaya IP Phones using Enterprise Device Manager

This chapter provides procedures you can use to configure Auto-Detection and Auto-Configuration (ADAC) using Enterprise Device Manager (EDM).

Navigation

Configuring ADAC globally using EDM

Use the following procedure to configure global ADAC settings for the switch.

Procedure steps

1. From the navigation tree, double-click **Edit** to open the Edit navigation tree.
2. Double-click **ADAC** to open the ADAC work area.
3. Click the **ADAC** tab.
4. Select the **AdminEnable** box to enable ADAC globally.

   OR

   Clear the **AdminEnable** box to disable ADAC globally.
5. Click an **OperatingMode** radio button.
6. Select the **NotificationControlEnable** check box to enable trap notifications globally.

OR

Clear the **NotificationControlEnable** check box to disable trap notifications globally.

7. In the **VoiceVlan** dialog box, type a value.
8. Click the **CallServerPort** elipsis (...).
9. From the Call Server port list, Select Call Server ports.
10. Click **Ok**.
11. Click the **UplinkPort** elipsis (...).
12. From the uplink port list, select uplink ports.
13. Click **Ok**.
14. Click a **MacAddrRangeControl** radio button.
15. Click **Apply**.

Important:

You cannot apply the global ADAC configuration if VoiceVLAN, CallServerPort, or UplinkPort boxes are set to 0 or empty when AdminEnable is selected and the operating mode is tagged frames or advanced untagged frames.

Important:

You cannot configure the same port values for Call Server and Uplink.

Variable definitions

Use the data in the following table to configure global ADAC .

Variable	Value
AdminEnable	Enables and disables ADAC.
OperEnable	Indicates ADAC operational state: true is enabled and false is disabled. **Important:** If AdminEnable is True and OperEnable is False, this indicates an error condition such as missing Uplink and Call Server ports.
OperatingMode	Configures the ADAC operation mode:

Variable	Value
	• untaggedFramesBasic: IP Phones send untagged frames, and the Voice VLAN is not created. • untaggedFramesAdvanced: IP Phones send untagged frames, and the Voice VLAN is created. • taggedFrames: IP Phones send tagged frames.
NotificationControlEnable	Enables or disables ADAC trap notifications.
VoiceVlan	Selects the Voice VLAN ID.
CallServerPort	Selects the Call Server port. A maximum of 8 Call Server ports are supported.
UplinkPort	Selects the Uplink port. A maximum of 8 uplink ports are supported.
MacAddrRangeControl	Provides two options for configuring the MAC address range table: • none: no MAC address range table selected. • clearTable: clears the MAC address range table. • defaultTable: sets the MAC address range table to its default values.

ADAC port configuration using EDM

Use the information in this section to configure ADAC for switch ports and to display port-based ADAC information.

Viewing port ADAC for informaion using EDM

Use the following procedure to display ADAC configuration information for switch ports.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. Double-click **Chassis** then double-click **Ports**.

   OR

Double-click **ADAC** .

3. In the Ports work area, click the **ADAC** tab.

   OR

   In the ADAC work area, click the **ADAC Ports** tab.

4. On the toolbar, you can click **Refresh** to update the data.

Variable definitions

Use the data in the following table to help you understand the ADAC display for switch ports.

Variable	Value
Index	Indicates the switch position in a stack and the port number. The default value for a standalone switch is 1.
AdminEnable	Indicates whether ADAC is enabled (true) or disabled (false) for the port.
OperEnable	Indicates ADAC operational state: true (enabled) or false (disabled). **Important:** If OperEnable is False and AdminEnable is True, then Auto-Detection/Auto-Configuration is disabled. This can occur due to a condition such as reaching the maximum number of devices supported per port.
ConfigStatus	Indicates the ADAC status for the port. Values include: • configApplied—the ADAC configuration is applied to this port. • configNotApplied—the ADAC configuration is not applied to this port.
TaggedFramesPvid	Indicates a unique PVID between 1 and 4094. A value of 0 indicates that Auto-Configuration cannot change the PVID for the port.
TaggedFramesTagging	Indicates the ADAC operating mode. Values include: • tagAll—tags all frames • tagPvidOnly—tags frames by the unique PVID • untagPvidOnly—untags frames by the unique PVID • noChange—accepts frames without change
AdacPortType	Indicates how ADAC classifies the port. Values include:

Variable	Value
	• telephony (when Auto-Detection is enabled for the port) • telephony—auto-detection is enabled • callServer —port is configured as a call server • uplink—port is configured as an uplink or is part of the same trunk as the uplink port • other—the port is not classified as either telephony, callServer, or uplink.
MacDetectionEnable	Indicates whether Auto-Detection of Avaya IP Phones, based on MAC address, is enabled (true) or disabled (false) on the interface.
LldpDetectionEnable	Indicates whether Auto-Detection of Avaya IP Phones, based on 802.1AB, is enabled (true) or disabled (false) on the interface. When cleared, indicates that Auto-Detection of Avaya IP Phones, based on 802.1AB, is disabled on the interface.

Configuring ADAC for specific ports using EDM

Use the following procedure to configure ADAC for one or more ports in a standalone switch or switch stack.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. Double-click **Chassis** then double-click **Ports**.

   OR

   Double-click **ADAC** .
3. In the Ports work area, click the **ADAC** tab.

   OR

   In the ADAC work area, click the **ADAC Ports** tab.
4. To select a port to edit, click the port **Index**.
5. In the port row, double-click the cell in the **AdminEnable** column.
6. Select a value from the list—true to enable ADAC for the port, or false to disable ADAC for the port.

7. In the port row, double-click the cell in the **TaggedFramesPvid** column.
8. Type a value in the dialog box.
9. In the port row, double-click the cell in the **TaggedFramesTagging** column.
10. Select a value from the list.
11. In the port row, double-click the cell in the **MacDetectionEnable** column.
12. Select a value from the list—true to enable MAC address detection for the port, or false to disable MAC address detection for the port.
13. In the port row, double-click the cell in the **LldpDetectionEnable** column.
14. Select a value from the list—true to enable LLDP detection for the port, or false to disable LLDP detection for the port.
15. Repeat steps **4** through **14** to configure ADAC for additional ports.
16. Click **Apply**.

Variable definitions

Use the data in the following table to configure ADAC for one or more switch ports.

Variable	Value
Index	Indicates the switch position in a stack and the port number. The default value for a standalone switch is 1.
AdminEnable	Enables or disables ADAC for the port.
OperEnable	Indicates ADAC operational state: true (enabled) or false (disabled). This is a read-only cell. **Important:** If OperEnable is False and AdminEnable is True, then Auto-Detection/Auto-Configuration is disabled. This can occur due to a condition such as reaching the maximum number of devices supported per port.
ConfigStatus	Indicates the ADAC status for the port. This is a read-only cell. Values include: • configApplied—the ADAC configuration is applied to this port. • configNotApplied—the ADAC configuration is not applied to this port.
TaggedFramesPvid	Specifies a unique PVID between 1 and 4094. A value of 0 indicates that Auto-Configuration cannot change the PVID for the port.

Variable	Value
TaggedFramesTagging	Specifies the ADAC operating mode. Values include: • tagAll—tags all frames • tagPvidOnly—tags frames by the unique PVID • untagPvidOnly—untags frames by the unique PVID • noChange—accepts frames without change
AdacPortType	Indicates how ADAC classifies the port. This is a read-only cell. Values include: • telephony (when Auto-Detection is enabled for the port) • telephony—auto-detection is enabled • callServer —port is configured as a call server • uplink—port is configured as an uplink or is part of the same trunk as the uplink port • other—the port is not classified as either telephony, callServer, or uplink.
MacDetectionEnable	When selected, indicates that Auto-Detection of Avaya IP Phones, based on MAC address, is enabled on the interface. When cleared, indicates that Auto-Detection of Avaya IP Phones, based on MAC address, is disabled on the interface. **Important:** MacDetectionEnable cannot be set to false if no other supported detection mechanism is enabled on the port.
LldpDetectionEnable	When selected, indicates that Auto-Detection of Avaya IP Phones, based on 802.1AB is enabled on the interface. When cleared, indicates that Auto-Detection of Avaya IP Phones, based on 802.1AB, is disabled on the interface. **Important:** LldpDetectionEnable cannot be set to False if no other supported detection mechanism is enabled on the port.

ADAC MAC address range configuration using EDM

Use the information in this section to manage the ADAC MAC address range table.

Viewing the MAC address range table using EDM

Use the following procedure to display the MAC address range table.

Procedure steps

1. From the navigation tree, click **Edit** to open the Edit navigation tree.
2. Double-click **ADAC** to open the Chassis work area.
3. Select the **ADAC MAC Ranges** tab.

Variable definitions

Use the data in the following table to help you understand the MAC address range table display.

Variable	Value
MacAddrRangeLowEndIndex	Indicates the low-end MAC address of the range.
MacAddrRangeHighEndIndex	Indicates the high-end MAC address of the range.

Creating MAC address ranges using EDM

Use the following procedure to add new MAC address ranges to the ADAC MAC address range table.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **ADAC**.
3. Click the **ADAC MAC Ranges** tab.
4. Click **Insert**.
5. In the **MacAddrRangeLowEndIndex** box, type the MAC address for the low end of the IP Phone MAC address range.

6. In the **MacAddrRangeHighEndIndex** box, type the MAC address for the high end of the IP Phone MAC address range.
7. Click **Insert**.
8. Click **Apply**.

Deleting MAC address ranges using EDM

Use the following procedure for deleting MAC address ranges to remove MAC address ranges from the ADAC MAC address range table.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **ADAC**.
3. Click the **ADAC MAC Ranges** tab.
4. Click the MAC address range to delete.
5. Click **Delete**.
6. Click **Yes** to confirm the deletion of the MAC address range from the table.

Appendix A: Quick configuration for Multi-Link Trunking

If you are a system administrator with experience configuring Ethernet Routing Switch 2500 Series Multi-Link Trunking, use the flowchart in Figure 18: Configuring Multi-Link Trunks on page 223 as a quick configuration guide. The flowchart refers you to the "configuration rules" appropriate for this feature.

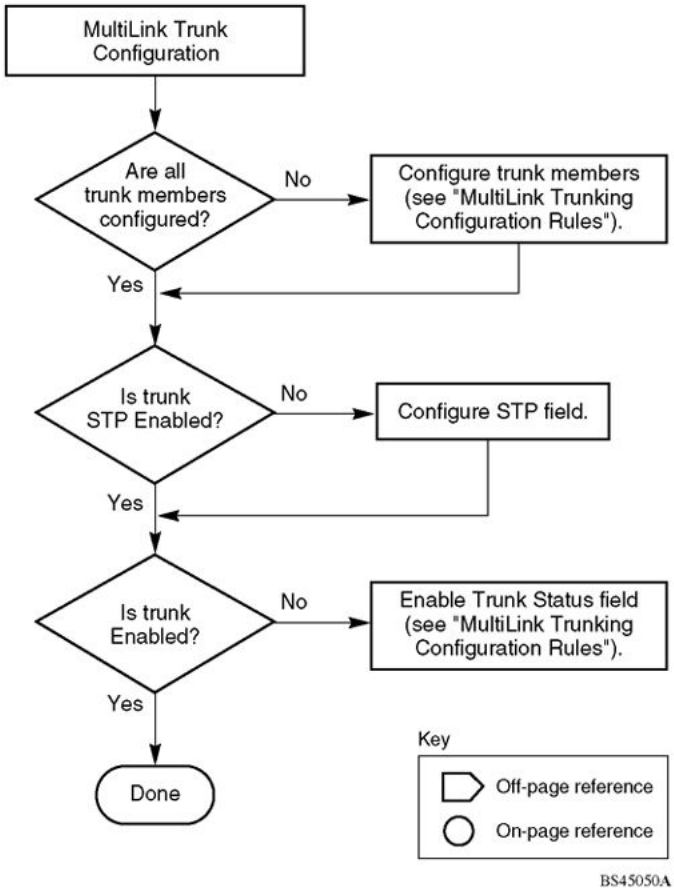

Figure 18: Configuring Multi-Link Trunks