



Configuration — Link Aggregation, MLT, and SMLT

Avaya Ethernet Routing Switch 8800/8600

7.1.3
NN46205-518, 06.01
January 2012

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/](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>.

Preventing Toll Fraud

"Toll fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud Intervention

If you suspect that you are being victimized by Toll Fraud and you need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support Web site: <http://support.avaya.com>. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

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: Purpose of this document.....	7
Chapter 2: New in this release.....	9
Features.....	9
VLACP flap detect and damping.....	9
Dynamic VLAN addition to LACP-enabled MLT (LAG).....	9
Other changes.....	9
Chapter 3: Link aggregation fundamentals.....	11
Link aggregation overview.....	11
MultiLink Trunking.....	12
Navigation.....	12
MLT traffic distribution algorithm.....	12
Multicast flow distribution over MLT.....	14
MLT and MLT with LACP configuration rules.....	18
LAG rules.....	19
MultiLink Trunking with LACP.....	20
IEEE 802.3ad overview.....	20
802.3ad link aggregation principles.....	22
Split MultiLink Trunking.....	23
Navigation.....	23
SMLT overview.....	24
SMLT versus STP.....	25
SMLT topologies.....	26
SMLT and Interswitch trunking.....	26
SMLT and IST traffic flow example.....	29
SMLT Interoperability with MSTP.....	31
Single Port SMLT.....	33
MLT-based SMLT with Single Port SMLT.....	35
SMLT and LACP support.....	37
SMLT and IP routing.....	38
SMLT and SLPP.....	39
SMLT/RSMLT operational improvements.....	40
Additional SMLT considerations.....	41
Virtual Link Aggregation Control Protocol.....	43
VLACP timers and sub-100-ms core convergence.....	45
VLACP considerations.....	46
VLACP flap detect and damping.....	47
LACP and Minimum Link.....	48
MinLink requirements.....	48
MinLink operations.....	49
MinLink considerations.....	49
Link aggregation configuration considerations.....	49
Navigation.....	49
MLT with LACP configuration considerations.....	50
MLT with LACP and SMLT configuration considerations.....	51

MLT with LACP and Spanning Tree configuration considerations.....	53
LACP parameters configuration considerations.....	54
Link aggregation and scaling.....	55
Chapter 4: Link aggregation configuration using Enterprise Device Manager.....	57
Configuring global LACP parameters.....	57
Configuring LACP parameters.....	59
Configuring LACP on a port.....	61
Chapter 5: MLT and SMLT link aggregation configuration using Enterprise Device Manager.....	67
Adding a MultiLink/LACP trunk.....	67
Adding ports to an MLT.....	70
Adding an MLT-based SMLT.....	70
Viewing MLT-based SMLTs.....	72
Adding ports to an MLT-based SMLT.....	73
Configuring an IST MLT.....	74
Editing an IST.....	75
Enabling SMLT Interoperability with MSTP using EDM.....	76
Disabling SMLT Interoperability with MSTP using EDM.....	76
Displaying IST MLT information using EDM.....	77
Configuring a single port SMLT.....	77
Viewing Single Port SMLTs.....	79
Deleting a Single Port SMLT.....	80
Chapter 6: VLACP configuration using Enterprise Device Manager.....	81
Enabling VLACP globally.....	81
Configuring VLACP on a port.....	81
Configuring VLACP flap detect and damping.....	83
Chapter 7: Link aggregation configuration using the CLI.....	85
Prerequisites.....	85
Job aid: Roadmap of LACP commands.....	85
Configuring global LACP parameters.....	87
Configuring LACP on a port.....	88
LACP show commands.....	91
Navigation.....	91
Viewing global LACP configuration information.....	91
Viewing LACP configuration information.....	91
Chapter 8: MLT and SMLT link aggregation configuration using the CLI.....	93
Job aid: Roadmap of MLT and SMLT commands.....	93
Configuring an MLT.....	95
Adding ports to an MLT LAG.....	96
Removing ports from an MLT LAG.....	97
Viewing MLT LACP configuration information for aggregators.....	98
Configuring LACP on an MLT.....	99
Creating a SMLT from an existing MLT.....	100
Creating an interswitch trunk.....	101
Creating an interswitch trunk from an existing MLT.....	102
Enabling and disabling an interswitch trunk.....	103
Deleting an interswitch trunk.....	104

Enabling SMLT Interoperability with MSTP using CLI.....	104
Disabling SMLT Interoperability with MSTP using CLI.....	105
Displaying MLT IST information using CLI.....	105
Configuring a single port SMLT.....	106
Viewing the MLT port calculated by the MLT hash algorithm.....	107
MLT and SMLT show commands.....	108
MLT and SMLT show commands navigation.....	108
Viewing MLT information.....	108
Viewing information about collision errors.....	109
Viewing information about Ethernet errors.....	109
Viewing MLT status.....	110
Viewing interswitch trunk status.....	110
Viewing SMLT status.....	111
Viewing all ports configured for single port SMLT.....	111
Viewing a port configured for Single Port SMLT.....	111
Chapter 9: VLACP configuration using the CLI.....	113
Job aid: Roadmap of VLACP commands.....	113
Configuring VLACP on a port.....	114
Viewing VLACP port configuration.....	115
Enabling or disabling VLACP.....	115
Configuring VLACP flap detect and damping.....	116
Viewing VLACP flap detect and damping configuration.....	117
Chapter 10: MLT and MLT with LACP CLI configuration examples.....	119
MultiLink Trunking configuration example.....	119
Multilink with LACP configuration example.....	120
Chapter 11: Link aggregation configuration using the ACLI.....	123
Prerequisites.....	123
Job aid: Roadmap of LACP commands.....	123
Configuring global LACP parameters.....	125
Configuring LACP on a port.....	127
Viewing LACP configuration information.....	129
Chapter 12: MLT and SMLT link aggregation configuration using the ACLI.....	131
Job aid: Roadmap of MLT and SMLT commands.....	131
Configuring an MLT.....	132
Adding ports to an MLT LAG.....	133
Removing ports from an MLT LAG.....	133
Configuring LACP on an MLT.....	134
Creating a SMLT from an existing MLT.....	135
Creating an interswitch trunk.....	136
Enabling an interswitch trunk.....	137
Deleting an interswitch trunk.....	137
Configuring a single port SMLT.....	138
Enabling SMLT Interoperability with MSTP using ACLI.....	139
Disabling SMLT Interoperability with MSTP using ACLI.....	139
Displaying MLT IST information using ACLI.....	140
Viewing the MLT port calculated by the MLT hash algorithm.....	140
MLT and SMLT show commands.....	141

MLT and SMLT show commands navigation.....	142
Viewing information about collision errors.....	142
Viewing information about Ethernet errors.....	143
Viewing all ports configured for single port SMLT.....	143
Chapter 13: VLACP configuration using the ACLI.....	145
Job aid: Roadmap of VLACP commands.....	145
Configuring VLACP on a port.....	146
Viewing the VLACP port configuration.....	147
Enabling or disabling VLACP globally.....	148
Configuring VLACP flap detect and damping.....	149
Viewing VLACP flap detect and damping configuration.....	150
Chapter 14: Customer Service.....	153
Getting technical documentation.....	153
Getting product training.....	153
Getting help from a distributor or reseller.....	153
Getting technical support from the Avaya Web site.....	153
Glossary.....	155

Chapter 1: Purpose of this document

This document contains procedural and conceptual information to help you configure and manage link aggregation and MultiLink trunking on the Avaya Ethernet Routing Switch 8800/8600. This document also provides instructions for using the Command Line Interface (CLI), the Avaya Command Line Interface (ACLI), and the Enterprise Device Manager (EDM).

For definitions of terms and acronyms used in this document, see *Avaya Ethernet Routing Switch 8800/8600 Terminology, NN46205-102*.

Purpose of this document

Chapter 2: New in this release

See the following sections for the changes in *Avaya Ethernet Routing Switch 8800/8600 Configuration — Link Aggregation, MLT, and SMLT*, NN46205-518 for Release 7.1.3.

- [Features](#) on page 9
- [Other changes](#) on page 9

Features

See the following sections for feature updates to this document in release 7.1.3.

VLACP flap detect and damping

VLACP flap detect and damping is used to automatically shut down selected VLACP links that are flapping due to link instability.

For more information, see

[VLACP flap detect and damping](#) on page 47

Dynamic VLAN addition to LACP-enabled MLT (LAG)

Starting in Release 7.1.3, you can add or remove a VLAN from an active LACP-enabled MLT without first disabling LACP on the MLT. For more information, see [MLT with LACP configuration considerations](#) on page 50

Other changes

There are no other changes in this release.

New in this release

Chapter 3: Link aggregation fundamentals

This chapter describes link aggregation concepts and the features supported on the Avaya Ethernet Routing Switch 8800/8600.

Link aggregation overview

Link aggregation provides link level redundancy and increases load sharing. Link aggregation allows you to bundle the ports into a port group, which is represented as one logical interface to the MAC layer.

The Avaya Ethernet Routing Switch 8800/8600 supports the following types of link aggregation:

- MultiLink Trunking (MLT)—a statically configured link bundling method. MLT is not standards based, but will interoperate with other vendors' static link methods.
- IEEE 802.3ad-based link aggregation, through the Link Aggregation Control Protocol (LACP), supports a dynamic link aggregation function as they become available to a trunk group. LACP dynamically detects whether links can be aggregated into a link aggregation group (LAG) and does so when links become available. LACP also provides link integrity checking at Layer 2 for all links within the LAG.

Both MLT and IEEE 802.3ad-based link aggregation are defined as point-to-point functions.

The Avaya Ethernet Routing Switch 8800/8600 software offers LACP functionality layered with MLT. This document uses the term MLT with LACP to refer to this functionality.

Split MultiLink Trunking (SMLT) is an option that improves Layer 2 (bridged) resiliency by providing for the addition of switch failure redundancy with sub-second failover, on top of all standard MLT link failure protection and flexible bandwidth scaling functionality. SMLT allows you to connect any device which supports some form of link aggregation, be it a switch or a server, to two distinct separate SMLT endpoints or switches. These SMLT switches form a Switch Cluster and are referred to as an IST Core Switch pair.

LACP can also be used on SMLT configurations. The Avaya Ethernet Routing Switch 8800/8600 provides modifications to LACP in SMLT configurations. This allows LACP-capable devices to connect to an SMLT aggregation pair. Avaya recommends that LACP not be configured on the IST MLT.

Virtual LACP (VLACP) is an Avaya modification that provides end-to-end failure detection. VLACP is not a link aggregation protocol; VLACP implements link status control protocol at the port level. It is a mechanism to periodically check the end-to-end health of a point-to-point or end-to-end connection. You can run VLACP on single ports or on ports that are part of a MLT. Avaya recommends that you do not configure VLACP on LACP-enabled ports. VLACP

does not operate properly with LACP. You can configure VLACP with any SMLT configuration.

MultiLink Trunking

MultiLink Trunking (MLT) is a point-to-point connection that aggregates multiple ports to logically act like a single port, with the aggregated bandwidth. Grouping multiple ports into a logical link provides a higher aggregate on a switch-to-switch or switch-to-server application.

MLT provides module redundancy via Distributed MultiLink Trunking (DMLT). DMLT allows you to aggregate similar ports from different modules. Avaya recommends always using DMLT when possible.

To include ports as trunk group members of an MLT, you must statically configure the ports.

Navigation

- [MLT traffic distribution algorithm](#) on page 12
- [Multicast flow distribution over MLT](#) on page 14
- [MLT and MLT with LACP configuration rules](#) on page 18
- [LAG rules](#) on page 19
- [MLT network topology and configuration examples](#) on page 16

MLT traffic distribution algorithm

You can use a MultiLink trunk to aggregate bandwidth between two switches. The Ethernet Routing Switch 8800/8600 uses one of two algorithms to determine which active port in the MultiLink trunk to use for each packet. The MLT algorithms provide load sharing, not load balancing. The MLT algorithms ensure that each packet in a flow does not arrive out of sequence, and that any one flow always traverses the same link path.

The algorithms are the same traffic distribution algorithms used for the IEEE 802.3ad-based link aggregation. Note that the distribution is determined by the ingress port.

To view port and algorithm calculations, use the procedure given in [Viewing the MLT port calculated by the MLT hash algorithm](#) on page 107.

You can use the ping snoop command as a tool to determine which link in an MLT a flow is traversing on.

R, RS, and 8800 module traffic distribution algorithm

This release uses an enhanced traffic distribution algorithm for R, RS, and 8800 modules, which ensures proper traffic distribution in all customer networks. The MLT hashing algorithm is based on Layer 4 fields for IPv4 UDP/TCP traffic for R, RS, and 8800 modules, versus Source and Destination IP addresses. This hash change provides a better distribution of traffic for many environments, with no impact on module performance. MLT hashing for any module type is determined at the ingress port, although the affect is seen at the egress MLT. That is, the traffic ingressing a port determines the destination is an MLT and sends the traffic to the correctly hashed MLT port, regardless of MLT port type or even a mixed port type. Therefore in a mixed chassis, the algorithm used is based on ingress port for the traffic, not the MLT configuration. The new hash algorithm for IPv4 TCP/UDP traffic is as follows:

64-bit key = (SrcPort (16 bits), DstPORT (16 bits), DstIP (LSB 16 bits), SrcIP (LSB 16 bits))

For non TCP/UDP IPv4 traffic:

64-bit key = (DstIP (32 bits), SrcIP (32 bits))

The 64-bit key result is used as an index to a table populated with active MLT ports repeated over 63 entries.

The CLI command **config sys set hash-calc getmltindex** is modified to allow src-port and dst-port as optional parameters. This parameter will provide the transmit or egress MLT index or port for ingress traffic.

Note that the new hashing for IP traffic between a given source and destination IP address will be different for TCP/UDP packets and ICMP packets. Therefore, you cannot use the ping-snoop feature to reliably determine the hashed path taken by IP TCP/UDP traffic when performed on an ingress port.

The enhanced hash algorithm does not apply to IPv6 traffic, just IPv4. The distribution algorithm for IPv6 traffic (bridged or routed) is:

- 64-bit key = (SrcIP[63:0] XOR srcIplp[127:64])
- 64-bit key = ((64-bit key) XOR DestIp[63:0])
- 64-bit key = ((64-bit key) XOR DestIp[127:64])

The distribution algorithm enhancement for other types of IPv6 traffic is to form a 64-bit hash key using the lower 32 bits of the DestMAC field and the lower 32 bits of the SrcMAC field:

64-bit key = hash(64-bit key, hashFcn)

where hash is a Remote Service Provider (RSP) instruction and hashFcn produces a 6-bit result. The hash function is 0x000c00003f000000, which contains a 25-bit seed, a 24-bit hash polynomial coefficient, and a 5-bit polynomial width.

The 64-bit key result is used as an index to a table populated with active MLT ports repeated over 63 entries.

MultiLink trunking and auto-negotiation interaction

In this release, MLT requires that all ports operate at the same speed. When you use MLT with LACP, LACP dynamically checks for proper speed on all port members. You are no longer required to have similar physical connection types. For example, you can mix a fiber port with a copper port. When you use auto-negotiation with MLT and not LACP, you need to ensure that all ports run at the same speed.

Multicast flow distribution over MLT

MLT provides a mechanism to distribute multicast streams over a MultiLink trunk. This mechanism is based on source-subnet and group addresses. You can use it to choose the address and bytes in the address for the distribution algorithm.

The multicast flow distribution over MLT algorithm is the same multicast flow distribution algorithm used in IEEE 802.3ad-based link aggregation. As a result, you can distribute the load on the MLT ports and achieve even stream distribution. In applications such as Television (TV) distribution, multicast traffic distribution is important because bandwidth requirements are substantial when a large number of TV streams are employed.

When you configure flow distribution over MLT, Avaya recommends that you choose source and group masks that result in even traffic distribution over the MultiLink trunk links. For example, if the group addresses change incrementally while few sources sending to different groups, use a source mask of 0.0.0.0 and a group mask of 255.255.255.255. In most cases, this provides a sequential distribution of traffic on the MultiLink trunk links.

When changing the state parameters of multicast flow distribution, minor traffic interruptions can occur.

For a detailed description of commands used to configure multicast flow distribution over MLT, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Multicast Routing Protocols, NN46205-501*.

Multicast distribution algorithm

To determine the port for a particular source, group (S, G) pair, use the number of active MLT ports to MOD. The number generated by the XOR (exclusive OR operation) for each byte of the masked group address, with the masked source address.

By default, the group mask and source mask is 255.255.255.255. A byte with a value of 255 in the mask indicates that the corresponding byte in the group or source address is taken into account when the algorithm is applied.

The following is an example:

group address G[0].G[1].G[2].G[3]

group mask GM[0].GM[1].GM[2].GM[3]

source subnet address $S[0].S[1].S[2].S[3]$

source mask $SM[0].SM[1].SM[2].SM[3]$

The port is calculated by using:

$$(((((G[0] \text{ AND } GM[0]) \text{ XOR } (S[0] \text{ AND } SM[0])) \text{ XOR } ((G[1] \text{ AND } GM[1]) \text{ XOR } (S[1] \text{ AND } SM[1]))) \text{ XOR } ((G[2] \text{ AND } GM[2]) \text{ XOR } (S[2] \text{ AND } SM[2]))) \text{ XOR } ((G[3] \text{ AND } GM[3]) \text{ XOR } (S[3] \text{ AND } SM[3]))) \text{ MOD (active ports of the MLT)}$$

Algorithm example

The traffic distribution algorithm causes sequential distribution if the streams are similar in the following example.

Assume that the MLT ports are 1/1 to 1/4, the source mask configuration is 0.0.0.0 and the group mask configuration is 255.255.255.255. The source A.B.C.D. is sent to the following groups:

X.Y.Z.1

X.Y.Z.2

X.Y.Z.3 to X.Y.Z.10

The algorithm chooses link 1/1 for group X.Y.Z.1, X.Y.Z.2 goes on port 1/2, X.Y.Z.3 goes on port 1/3, X.Y.Z.4 goes on port 1/4, X.Y.Z.5 goes on port 1/1, and so on.

Multicast traffic distribution

The goal of traffic distribution is to achieve stream distribution on the MultiLink trunk links when an MLT configuration change takes place. Traffic distribution is based on the IP multicast MLT distribution algorithm, which is the same for all modules.

Traffic distribution is disabled by default. When you add or remove a link from the MultiLink trunk, the active streams continue to flow on the original links. If distribution is enabled, the active streams redistribute according to the distribution algorithm on the links of the MultiLink trunk links.

When changing the state parameters of multicast traffic distribution, minor traffic interruptions can occur.

To minimize the effect of multicast traffic distribution on the MultiLink trunks, the implementation does not move the streams to the appropriate links at one time. Instead, it redistributes a few streams at every time tick of the system. When an MLT port becomes inactive and distribution is disabled, only affected streams are redistributed on the remaining active ports.

If you enable distribution, all streams redistribute on the MLT ports based on the distribution algorithm assignment.

If you disable distribution and a new port becomes active in a Multilink trunk, existing streams remain on the original links. Enable distribution to redistribute the streams dynamically and split the load on all the links of the MultiLink trunk. This redistributes a few streams at every system time tick.

For a detailed description about the commands you use to configure multicast flow distribution over MLT, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Multicast Routing Protocols, NN46205-501*.

MLT network topology and configuration examples

The following reference information contains examples of MLT network topology and configuration. The same topologies apply to MLT with LACP.

Example 1: Switch-to-switch MLT

[Figure 1: Switch-to-switch MultiLink trunks configuration](#) on page 16 shows two MultiLink trunks (T1 and T2) connecting switch S1 to switches S2 and S3.




Figure 1: Switch-to-switch MultiLink trunks configuration

In this example, you can configure each trunk with multiple switch ports to increase bandwidth and redundancy. When traffic between switch-to-switch connections approaches single port bandwidth limitations, you can create an MultiLink trunk to supply the additional bandwidth required to improve performance, as well as providing physical link layer redundancy.

Example 2: Switch-to-server MLT

[Figure 2: Switch-to-server MultiLink trunk configuration](#) on page 17 shows a typical switch-to-server trunk configuration. In this example, file server FS1 uses dual MAC addresses, with one MAC address for each Network Interface Card (NIC). No MultiLink trunk is configured on FS1. FS2 is a single MAC server (with a four-port NIC) configured as MultiLink trunk configuration T1.

In this example, one port on FS1 is blocked and unused and FS2 benefits from aggregated bandwidth on MultiLink trunk T1.




Figure 2: Switch-to-server MultiLink trunk configuration

Example 3: Client/Server MLT

[Figure 3: Client/server MultiLink trunk configuration](#) on page 18 shows how you can use MultiLink trunks in a client/server configuration. In this example, both servers are connected directly to switch S1. FS2 is connected through a MultiLink trunk configuration (T1). The switch-to-switch connections are through MLT T2, T3, and T4. Clients access data from the servers (FS1 and FS2) and receive maximized bandwidth through T1, T2, T3, and T4. On the Avaya Ethernet Routing Switch 8800/8600, trunk members (the ports that comprise each MultiLink

trunk) are not consecutive switch ports so to avoid module redundancy select different modules.




Figure 3: Client/server MultiLink trunk configuration

MLT and MLT with LACP configuration rules

Avaya Ethernet Routing Switch 8800/8600 MultiLink trunks adhere to the following rules. The rules also apply to MLT with LACP.

- MLT is supported on 10Base-T, 100Base-TX, 100Base-FX, Gigabit Ethernet, and 10 Gigabit Ethernet module ports.
- When you create an MLT with more than one port that ends in an 8648GTR module and the MLT is connected to a 10/100 Mbps device with Autonegotiation enabled, the link state for the first port is 100 Mbps full duplex and the state of all other MLT ports is 10 Mbps half duplex. Adjust the speed and duplex modes manually on a per port basis.
- All MultiLink trunk ports must have the same speed and duplex settings, even when auto-negotiation is set.

- The media type of MLT ports can be different; a mix of copper and fiber are allowed.
- All MultiLink trunk ports must be in the same Spanning Tree Group (STG) unless the port is tagged. Tagging allows ports to belong to multiple STGs, as well as multiple VLANs.
- When the port is made a member of MLT, it inherits the properties of the MLT and hence the STG properties are inherited from the VLAN associated with that MLT. When the port gets removed from MLT or when the MLT is deleted, the ports are removed from the MLT STG and added into the default STG.
- MLT is compatible with Spanning Tree Protocol (STP), Multiple Spanning Tree Protocol (MSTP) (IEEE 802.1s), and Rapid Spanning Tree Protocol (RSTP) (IEEE 802.1w).
- Tagging (IEEE 802.1Q) is supported on a MultiLink trunk.

Ethernet Routing Switch 8800/8600 MultiLink trunks have the following general features and requirements:

- Supports up to 128 MLT groups with as many as 8 same-type ports belonging to a single MultiLink trunk.
- MLT ports can span modules, providing module redundancy.
- Apply filters individually to each port in a MultiLink trunk.

With the STP enabled, ports in the same MultiLink trunk operate as follows:

- Each port sends identical Bridge Protocol Data Unit (BPDU).
- The MultiLink trunk port ID is the ID of the lowest numbered port.
- If identical BPDUs are received on all ports, the MultiLink trunk mode is forwarding.

You can disable the Avaya STP (**ntstg <enable|disable>**) if you do not want to receive BPDUs on all ports.

- If ports do not receive BPDUs on a port or BPDU and port tagging do not match, the individual port is taken offline.
- Path cost is inversely proportional to the active MultiLink trunk bandwidth.

LAG rules

The Ethernet Routing Switch 8800/8600 link aggregation group (LAG) adhere to the following rules:

- All LAG ports operate in full-duplex mode.
- All LAG ports operate at the same data rate.
- All LAG ports should belong to the same set of VLANs.
- Link aggregation is compatible with the STP, MSTP, and RSTP.
- Assign all LAG ports to the same STP groups.

- Ports in a LAG can exist on different modules.
- Support a maximum of 128 link aggregation groups.
- A LAG can be configured with up to 16 ports (8 active and 8 standby ports). Even if you have less than 16 ports configured, you still cannot have an aggregation group with more than 8 active (or 8 standby) ports at any one time.
- After a MultiLink trunk is configured with LACP, you cannot add or delete ports manually without first disabling LACP.

MultiLink Trunking with LACP

MultiLink Trunking (MLT) with Link Aggregation Control Protocol (LACP) manages switch ports and port memberships to form a link aggregation group (LAG). LACP allows you to gather one or more links to form a LAG, which a Media Access Control (MAC) client treats as a single link. LACP can dynamically add or remove LAG ports, depending on availability and state.

IEEE 802.3ad overview

The IEEE 802.3ad standard comprises service interfaces, the LACP, the Marker Protocol, link aggregation selection logic, a parser/multiplexer, frame distribution, and frame collection functions.

[Figure 4: Link aggregation sublayer \(according to IEEE 802.3ad\)](#) on page 21 shows the major functions of IEEE 802.3ad defined as multiple link aggregation.




Figure 4: Link aggregation sublayer (according to IEEE 802.3ad)

The link aggregation sublayer is composed of the following functions:

- **Frame distribution:**

This block takes frames submitted by the MAC client and sends them for transmission on the appropriate port based on a frame distribution algorithm employed by the Frame Distributor.

Frame distribution also includes an optional Marker Generator/Receiver used for the Marker Protocol. Ethernet Routing Switch 8800/8600, only implements the Marker Receiver function. For details about the frame distribution algorithm, see [MLT traffic distribution algorithm](#) on page 12.

- **Frame collection:**

This block passes frames received from the various ports to the MAC client. Frame collection also includes a Marker Responder which is used for the Marker Protocol.

- **Aggregator Parser/Multiplexers:**

During transmission operations, these blocks pass frame transmission requests from the Distributor, Marker Generator, and Marker Responder to the appropriate port.

During receive operations, these blocks distinguish among Marker Request, Marker Response, MAC Client Protocol Data Units (PDUs), and pass the blocks to the appropriate entity (Marker Responder, Marker Receiver, and Collector, respectively).

- **Aggregator:**

The combination of frame distribution and collection, and aggregator Parser/Multiplexers.

- **Aggregation Control:**

This block configures and controls link aggregation. It incorporates LACP for the automatic communication of aggregation capabilities between systems and automatic configuration of link aggregation.

- **Control Parser/Multiplexers:**

During transmission operations, these blocks pass frame transmission requests from the aggregator and Control entities to the appropriate port.

During receive operations, these blocks distinguish Link Aggregation Control Protocol Data Units (LACPDUs) from other frames. The blocks pass, passing the LACPDUs to the appropriate sublayer entity and all other frames to the aggregator.

802.3ad link aggregation principles

Link aggregation allows you to group switch ports together to form a link group to another switch or server. Link groups increase aggregate throughput between devices and provide link redundancy.

Link aggregation employs the following principles and concepts:

- A MAC client communicates with a set of ports through an aggregator, which presents a standard IEEE 802.3 service interface to the MAC client. The aggregator binds to one or more ports within a system.
- The aggregator distributes frame transmissions from the MAC client to various ports, collects received frames from the ports, and transparently passes the frames to the MAC client.
- A system can contain multiple aggregators serving multiple MAC clients. A port binds to a single aggregator at a time. A MAC client is served by a single aggregator at a time.
- The Link Aggregation Control function binds ports to aggregators within a system. The control function aggregates links, binds the system ports to an appropriate aggregator, and monitors conditions to determine when a change in aggregation is needed. Network managers can manually provide link aggregation control by manipulating the link aggregation state variables (for example, keys). You can also use LACP to automatically determine, configure, bind, and monitor link aggregation.

- The LACP uses peer exchanges across links to continually determine the aggregation capability of the links and provide the maximum level of aggregation capability between a pair of systems.
- Frame ordering is maintained for certain sequences of frame exchanges between MAC Clients. The distributor ensures that all frames of a conversation pass to a single port. The collector passes frames to the MAC client in the order they are received from the port. The collector can select frames received from the aggregated ports in any order. Since the frames are not ordered on a single link, this guarantees that frame ordering is maintained for any conversation.
- Conversations move among ports within an aggregation for load balancing and for maintaining availability if a link fails.
- The standard does not impose any particular distribution algorithm on the distributor. The Avaya Ethernet Routing Switch 8800/8600 uses the MLT traffic distribution algorithm described in [MLT traffic distribution algorithm](#) on page 12.
- Each port is assigned a unique, globally administered MAC address.

When entities initiate frame exchanges within the link aggregation sublayer, the source address is the MAC address. An example of an entity that initiates frame exchanges is LACP and Marker Protocol exchanges.

- Each aggregator is assigned a unique, globally administered MAC address that is used from the perspective of the MAC client, both as a source address for transmitted frames and as the destination address for received frames. You can use one of the port MAC addresses in the associated LAG as the MAC address of the aggregator.

Split MultiLink Trunking

SMLT is an option that improves Layer 2 and Layer 3 resiliency. The following sections discuss SMLT in more detail.

Navigation

- [SMLT overview](#) on page 24
- [SMLT versus STP](#) on page 25
- [SMLT topologies](#) on page 26
- [SMLT and Interswitch trunking](#) on page 26
- [SMLT environment traffic flow rules](#) on page 28
- [SMLT and IST traffic flow example](#) on page 29
- [Single Port SMLT](#) on page 33

- [MLT-based SMLT with Single Port SMLT](#) on page 35
- [SMLT and LACP support](#) on page 37
- [SMLT and IP routing](#) on page 38
- [SMLT and SLPP](#) on page 39

SMLT overview

Split MultiLink Trunking (SMLT) is an option that improves Layer 2 (bridged) resiliency by providing for the addition of switch failure redundancy with sub-second failover, on top of all standard MLT link failure protection and flexible bandwidth scaling functionality. SMLT allows you to connect any device which supports some form of link aggregation, be it a switch or a server, to two distinct separate SMLT endpoints or switches. These SMLT switches form a Switch Cluster and are referred to as an IST Core Switch pair.

Switch Clusters are always formed as a pair, but pairs of clusters can be combined in either a square or full-mesh fashion to increase the size and port density of the Switch Cluster. When configured in a Layer 3 or routed topology, the configuration is referenced as Routed SMLT (RSMLT). For information about Routed SMLT, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing* (NN46205-523).



Important:

In the case of a full-mesh SMLT configuration between 2 Clusters running OSPF (more likely an RSMLT configuration) because of the way that MLTs work in regards to CP-generated traffic, Avaya highly recommends that the MLT port (or ports) that form the square leg of the mesh (versus the cross connect) be placed on lowered number slot/port than the cross connections. The reason for this is because CP-generated traffic is always sent out on the lowered numbered ports when active. Using this recommendation will keep some OSPF adjacency up if all the links of the IST fail. Otherwise the switches which have a failed IST could lose complete OSPF adjacency to both switches in the other Cluster and therefore become isolated.

SMLT connections can be formed via single links from the switch cluster to the edge connection, Single Link SMLT (SLT), or via standard MLTs, or MLTs with LACP. Optionally, SMLT links can have VLACP enabled as well. These various link connections can be mixed. Within the same Switch Cluster, both SMLT and RSMLT can be configured, allowing a mixture of both Layer 2 and Layer 3 VLANs. For examples of various SMLT configurations, refer to [SMLT topologies](#) on page 26.



Important:

In SLT configurations where the SLT port and all the IST_MLT ports are configured on the same physical module (not a best practice design; IST_MLT should always be a D-MLT), the switch can experience FDB entry learning issues for MAC addresses learned across the IST. This can lead to connectivity loss. This condition can occur if the module is physically swapped out during switch operation, or if the module is disabled/enabled via CLI or EDM. The switch relearns the FDB entries after the **fdb-ageout time**, or you can manually relearn them by performing an **fdb-entry flush** on the relevant VLANs.

SMLT networks do not need to use the IEEE 802.1d STP to enable loop-free triangle topologies because SMLT inherently avoids loops due to its superior enhanced link aggregation protocol.

This is accomplished by implementing a method that allows two aggregation switches to appear as a single device to edge switches, which dual-home to the aggregation switches. The aggregation switches interconnect using an interswitch trunk, which allows them to exchange addressing and state information (permitting rapid fault detection and forwarding path modification). SMLT is designed for Layer 2 network connectivity, but can be configured in Layer 3 networks by working with VRRP or RSMLT Layer 2 edge.

For more information about SMLT configuration, see *Technical Configuration Guide for Switch Clustering using Split-Multilink Trunking (SMLT)*, NN48500-518 and *SMLT and RSMLT Technical Configuration Guide*. You can find the technical configuration guides at www.avaya.com/support with the rest of the ERS 8800/8600 documentation.

SMLT advantages

SMLT eliminates all single points of failure and creates multiple paths from all user access switches to the network core. In case of failure, SMLT recovers as quickly as possible using all capacity. SMLT provides a transparent and interoperable solution that requires no modification on the part of the majority of existing user access devices.

SMLT improves the reliability of Layer 2 networks that operate between user access switches and the network center aggregation switch by providing:

- load sharing among all links
- fast failover in case of link failures
- elimination of single point of failure
- fast recovery in case of nodal failure
- transparent and interoperable solutions
- removal of STP convergence issues

SMLT versus STP

Networks designed to have user access switches dual-home to two aggregation switches, and have VLANs spanning two or more user access switches, experience the following design constraints:

- No load sharing exists over redundant links
- Network convergence is slow in case of failure

With the introduction of SMLT, all dual-home Layer 2 frame-switched network devices with dual homes are no longer dependent on the STP for loop detection. A properly designed SMLT network inherently does not have any logical loops.

SMLT solves the spanning tree problem by combining two aggregation switches into one logical MLT entity, thus making it transparent to any type of edge switch. In the process, it provides quick convergence, while load sharing across all available trunks.

SMLT topologies

There are three generic topologies in which you can deploy SMLT:

- A triangle topology
- A square topology
- A full-mesh topology (depending on the resiliency and redundancy required)

For information about SMLT topologies and what topologies are supported per product, see *Switch Clustering (SMLT/SLT/RSMLT/MSMLT) Supported Topologies and Interoperability with ERS 8600 / 5500 / 8300 / 1600, NN48500-555*.

SMLT and Interswitch trunking

[Figure 5: SMLT aggregation switches and operations](#) on page 27 illustrates an SMLT configuration with a pair of Avaya Ethernet Routing Switch 8800/8600 devices (E and F) as aggregation switches and four separate user access switches (A, B, C, and D).

You must connect SMLT aggregation switches through an interswitch trunk. For example, user access switches B and C connect to the aggregation switches through MultiLink trunks split between the two aggregation switches. As shown in [Figure 5: SMLT aggregation switches and operations](#) on page 27, the implementation of SMLT only requires two SMLT-capable aggregation switches. You must connect these switches through an interswitch trunk.




Figure 5: SMLT aggregation switches and operations

Aggregation switches use the interswitch trunk to:

- Confirm that they are alive and exchange MAC address forwarding tables.
- Send traffic between single switches attached to the aggregation switches.
- Serve as a backup if one SMLT link fails.

Because SMLT requires the interswitch trunk, Avaya recommends that you use multiple links on the interswitch trunk to ensure reliability and high availability. Avaya also recommends that you use Gigabit Ethernet links for interswitch trunk connectivity to provide enough bandwidth for potential cross traffic.

Avaya recommends that an interswitch MultiLink trunk contain at least two physical ports. Avaya also recommends that you disable CP-Limit on all physical ports that are members of an interswitch trunk MultiLink trunk, so as not to compromise the stability of SMLT.

! Important:

Avaya recommends that you do not configure CP-Limit or Extended CP-Limit on the IST MLT.

Disabling the CP-Limit on interswitch trunk MultiLink trunk ports forces another, less-critical port to be disabled if the defined CP Limits are exceeded. This preserves network stability if a protection condition (CP Limit) arises. It is likely that one of the MultiLink trunk ports (risers) can be disabled in such a condition; traffic continues to flow uninterrupted through the remaining SMLT ports.

For more information about CP-Limit and Extended CP-Limit, see *Avaya Ethernet Routing Switch 8800/8600 Administration, NN46205-605*.

Other SMLT aggregation switch connections

The example shown in [Figure 5: SMLT aggregation switches and operations](#) on page 27 includes end stations which connect to each of the switches. In this example, a, b1, b2, c1, c2, and d are clients and printers, while e and f are servers and routers.

User access switches B and C can use any method to determine which link of the MultiLink trunk connections to use to forward a packet, as long as the same link is used for a Source Address and Destination Address (SA/ DA) pair. This is true regardless of whether B or C knows the DA. SMLT aggregation switches always send traffic directly to a user access switch, and only use the interswitch trunk for traffic that they cannot forward in another, more direct way.

SMLT environment traffic flow rules

Traffic flow in an SMLT environment adheres to the following rules:

- If a packet is received from an interswitch trunk port, it is not forwarded to any active SMLT groups. This is key in preventing network loops.



Note:

If you want to decommission an SMLT switch, disable the ports first to prevent loops in the network.

- When a packet is received, a look-up is performed on the forwarding database. If an entry exists, and if the entry was learned locally from the split MultiLink trunk or through the interswitch trunk as a remote split MultiLink trunk, it is forwarded to the local port (the packet must not be sent to the interswitch trunk for forwarding unless there is no local connection). Unknown and Broadcast packets flood out all ports that are members of this VLAN.
- For load sharing purposes in an SMLT configuration, the Avaya Ethernet Routing Switch 8800/8600 obeys the MLT traffic distribution algorithm.

SMLT traffic flow examples

The following traffic flow examples are based on [Figure 5: SMLT aggregation switches and operations](#) on page 27.

Example 1: Traffic flow from a to b1 or b2

Assuming a and b1/ b2 are communicating through Layer 2, traffic flows from A to switch E and is forwarded over its direct link to B. Traffic coming from b1 or b2 to a is sent by B on one of its MultiLink trunk ports.

B can send traffic from b1 to a on the link to switch E, and traffic from b2 to a on the link to F. In the case of traffic from b1, switch E forwards the traffic directly to switch A, while traffic from b2, which arrived at F, is forwarded across the interswitch trunk to E and then to A.

Example 2: Traffic flow from b1/ b2 to c1/ c2

Traffic from b1/ b2 to c1/ c2 is always sent by switch B through its MultiLink trunk to the core. No matter at which switch (E or F) arrives at, it is sent directly to C through the local link.

Example 3: Traffic flow from a to d

Traffic from a to d (and d to a) is forwarded across the interswitch trunk because it is the shortest path. This is treated as a standard link; SMLT and interswitch trunk parameters are not considered.

Example 4: Traffic flow from f to c1/ c2

Traffic from f to c1/ c2 is sent out directly from F. Return traffic from c1/ c2 allows you to have one active VRRP Master for each IP subnet. It is passed across the interswitch trunk if switch C sends it to E.

SMLT and IST traffic flow example

In an SMLT environment, the two aggregation switches share the same forwarding database by exchanging forwarding entries using the IST. In [Figure 7: Output of the command show vlan info fdb-e 10](#) on page 31, the forwarding databases are shown for a pair of IST nodes (B and C). The entry for 00:E0:7B:B3:04:00 is shown on node C as being learned on MLT-1, but because SMLT REMOTE is true, this entry was actually learned from node B. On B, that same entry is shown as being directly learned through MLT-1 because SMLT REMOTE is false. [Figure 6: Network topology for traffic flow example](#) on page 30 shows the network topology.




Figure 6: Network topology for traffic flow example

When a packet arrives at node C destined for 00:E0:7B:B3:04:00, if the SMLT REMOTE status is true, the switch tries to send the packet to MLT-1 rather than through the interswitch trunk. Traffic rarely traverses the interswitch trunk unless there is a failure. If this same packet arrives at B, it is then forwarded to MLT-1 on the local ports.

Response from B:						
ERS-8610:5# show vlan info fdb-e 10						
=====						
Vlan Fdb						
=====						
VLAN	MAC			QOS	SMLT	
ID STATUS	ADDRESS	INTERFACE	MONITOR	LEVEL	REMOTE	

10 self	00:00:5e:00:01:01	-	false	1	false	
10 learned	00:80:2d:ba:d6:01	MLT-5	false	1	true	
10 self	00:80:2d:be:22:01	-	false	1	false	
10 learned	00:e0:7b:b3:04:00	MLT-1	false	1	false	
10 learned	00:e0:7b:b3:04:10	MLT-1	false	1	true	
Response from C:						
=====						
Vlan Fdb						
=====						
VLAN	MAC			QOS	SMLT	
ID STATUS	ADDRESS	INTERFACE	MONITOR	LEVEL	REMOTE	

10 self	00:00:5e:00:01:01	-	false	1	false	
10 self	00:80:2d:ba:d6:01	-	false	1	false	
10 learned	00:80:2d:be:22:01	MLT-5	false	1	true	
10 learned	00:e0:7b:b3:04:00	MLT-1	false	1	true	
10 learned	00:e0:7b:b3:04:10	MLT-1	false	1	false	

Figure 7: Output of the command show vlan info fdb-e 10

SMLT Interoperability with MSTP

Release 7.1 introduces functionality that allows interoperability between an SMLT cluster and MSTP in an Ethernet network. When you configure an InterSwitch Trunk (IST) and enable SMLT, with one command you can enable interoperability between the SMLT cluster and the MSTP side of the network.

SMLT Interoperability for MSTP is a cost-effective advantage when, for instance, you deploy SMLT in the core and data center but need fiber connections to edge devices. Instead, you can deploy MSTP at the edge. Then you can connect the SMLT side of the network and the MSTP side using SMLT Interoperability for MSTP, without sacrificing transition or reconvergence times.

You can also use SMLT Interoperability for MSTP to provide resiliency and redundancy. For example, when you are migrating your network from a Spanning Tree design towards an SMLT Clustering design you can extend Layer 2 VLANs between the old and new networks and maintain redundancy in case of link or node failure.

If you want to disable SMLT Interoperability for MSTP, ensure that you remove redundant paths, for VLANs, to the existing IST path. To prevent the formation of a bridging loop and

network failure, do not break or remove an MSTP device chain before you disable SMLT interoperability.

Assumptions:

- All bridges in the network have a bridge identifier greater than 00:00:00:00:00:00:02.

General limitations:

- The SMLT cluster always acts as the Root Bridge within the MSTP domain. For MSTP, the SMLT Cluster is the CIST Regional Root, CIST Root, and MSTI Root for all MSTI instances.
- Only 1 SMLT cluster can participate in the MSTP domain: if more than one cluster tries to participate each advertises the same Root Bridge ID and a conflict occurs.
- No link or MLT connected in parallel to the IST MLT is supported.
- SMLT and SLT links cannot run MSTP.
- Only MSTP modes are supported for SMLT Interoperability.
- You must enable SMLT Interoperability on both switches forming the SMLT cluster. Layer 2 loops form if you enable SMLT Interoperability on only 1 member of the SMLT cluster.

To ensure that one of the switches forming the SMLT Cluster is always the Root bridge, when you enable SMLT Interoperability the system automatically assigns MSTP Bridge IDs 00:00:00:00:00:00:01 and 00:00:00:00:00:00:02.

With MSTP, and based on VLAN membership, physical ports can belong to multiple MSTP instances (CIST + 0 or more MSTIs). When you enable SMLT Interoperability with MSTP, because all VLANs are q-tagged over the IST link, the system enables STP on the IST and sets the PathCost to 1 across all MSTP instances.

MSTP does not allocate a Bridge MAC for MSTI instances; a MAC is allocated only for the Base CIST. But the system can set Bridge Priority values for MSTI instances. So, for MSTP operation with SMLT Interoperability, the system changes the CIST Base MAC to 00:00:00:00:00:01 and 00:00:00:00:00:02, like the RSTP implementation, while the system sets the Bridge Priority to 00:00 across all instances (CIST & any MSTIs).

When you activate SMLT Interoperability with MSTP, the following list summarizes what happens automatically for all MSTP instances (CIST/MSTIs):

- STP is enabled on the IST ports for all instances.
- STP PathCost is set to 1 on IST ports for all instances.
- The STP Bridge Priority is set to 00:00 for all instances.
- The MSTP Base MAC is set to 00:00:00:00:00:01 on the switch with the highest IST IP (SMLT master).
- The SMLT master is always the CIST Root + Regional Root + MSTI Root.
- The RSTP Base MAC is set to 00:00:00:00:00:02 on the switch with the lowest IST IP (SMLT slave).
- The SMLT slave is always the CIST/Regional/MSTI Root if the SMLT master IST peer fails.

Limitations:

- You cannot directly configure or modify any MSTP settings (CIST or MSTI) on the IST ports.
- You cannot configure the Bridge Priority on any instance (CIST or MSTI), which is hard coded to 00:00.

You cannot stagger the Root bridge across the two SMLT peers, that is one switch cannot be the MSTP Root for the CIST + MSTI2 while the other switch is MSTP Root for MSTI1 + MSTI3. But you can still achieve load balancing over alternate active forwarding topologies by changing the PathCost setting on the MSTP nodes connected to your SMLT Cluster.

Single Port SMLT

Use Single Port SMLT to configure a split MultiLink trunk that uses a single port. The Single Port SMLT behaves like an MLT-based SMLT and can coexist with SMLTs in the same system. Single Port SMLT lets you scale the number of split MultiLink trunk on a switch to the maximum number of available ports.

Split MultiLink trunk links can exist in the following combinations on the SMLT aggregation switch pair:

- MLT-based SMLT and MLT-based SMLT
- MLT-based SMLT and single link SMLT
- single link SMLT and single link SMLT

Single Port SMLT configuration

You can use Single Port SMLT when you exceed the Avaya Ethernet Routing Switch 8800/8600 MultiLink trunk Group ID limit for server farm applications. The Single Port SMLT topology allows scaling up to the maximum number of ports on a switch. You can use any Layer 2 switch capable of link aggregation as the client in this case.




Figure 8: Single Port SMLT configuration at the port level

Single Port SMLT configuration considerations

The rules for configuring Single Port SMLT include:

- The dual-home device with dual homes that connect to the aggregation switches must support MLT.
- Single Port SMLT is supported on Ethernet ports.
- Each Single Port SMLT is assigned an SMLT ID from 1 to 512.
- You can designate Single Port SMLT ports as Access or Trunk (IEEE 802.1Q tagged or not); changing the type does not affect behavior.
- You cannot change a Single Port split MultiLink trunk to an MLT-based split MultiLink trunk by adding additional ports. You must delete the single port split MultiLink trunk and reconfigure the port as SMLT/MLT.
- You cannot change an MLT-based split MultiLink trunk into a single port split MultiLink trunk by deleting all ports except one. You must remove the SMLT/MLT and reconfigure the port as Single Port SMLT.

- You cannot configure a port as an MLT-based Split MultiLink trunk and as a single port SMLT at the same time.
- Two or more aggregation switches can have single port Split MultiLink trunk with the same IDs. You can have as many single port Split MultiLink trunk as there are available ports on the switch.
- LACP is supported on single port SMLT.

MLT-based SMLT with Single Port SMLT

You can configure an split MultiLink trunk with a single port split MultiLink trunk on one side and an MLT-based split MultiLink trunk on the other. Both must have the same SMLT ID.

Note:

You can configure Link Aggregation/MLT group IDs from 1 through 256. The switch supports a total of 128 MLT groups. Avaya suggest that users start numbering S-SMLT or SLT IDs to avoid overlapping MLT IDs in the 1 through 128 range. SLT IDs range from 1 through 512. If you start at 129, a switch can support the full capacity of 383 SLT ports or IDs.

In addition to general use, the [Figure 9: Changing a split trunk from MLT-based SMLT to Single Port SMLT](#) on page 36 shows how this configuration can be used to upgrade an MLT-based SMLT to a single port split MultiLink trunk without stopping the split trunk. Assuming the MLT ID is 10, the steps are:

1. Configure switch A and B with MTL-based SMLT.
2. Disable all ports and then delete MLT-based SMLT 10 on switch B. All traffic switched over to SMLT 10 on switch A.
3. Configure single port SMLT ID 10 on switch B and then enable the port. Traffic switches over both sides of the split trunk.
4. Disable all ports and then delete MLT-based SMLT 10 on switch A. All traffic switches over single port SMLT 10 on switch B.
5. Configure single port SMLT 10 on switch A and then enable the port. Traffic switched over both sides of split trunk.




Figure 9: Changing a split trunk from MLT-based SMLT to Single Port SMLT

SMLT and LACP support

The Ethernet Routing Switch 8800/8600 fully supports the IEEE 802.3ad LACP on MLTs and on a pair of SMLT switches.

With LACP the Ethernet Routing Switch 8800/8600 provides a standardized external link aggregation interface to third-party vendor IEEE 802.3ad implementations. This protocol extension provides dynamic link aggregation mechanisms. Only dual-home devices benefit from this enhancement.

Advantages of this protocol extension include:

- MLT peers and SMLT client devices can be both network switches and any type of server/workstation that supports link bundling through IEEE 802.3ad.
- Single-link and MultiLink trunk solutions support dual-home connectivity for more than 350 attached devices, so that you can build dual-home server farm solutions.

Supported SMLT/LACP scenarios

SMLT/IEEE link aggregation interaction supports all known SMLT scenarios in which an IEEE 802.3ad SMLT pair connects to SMLT clients, or in which two IEEE 802.3ad SMLT pairs connect to each other in a square or full-mesh topology.

Unsupported SMLT/LACP scenarios

Some of the unsupported SMLT/LACP scenarios include the following factors, which lead to failure:

- Incorrect port connections.
- Mismatched SMLT IDs assigned to SMLT client. SMLT switches can detect if SMLT IDs are not consistent. The SMLT aggregation switch, which has the lower IP address, does not allow the SMLT port to become a member of the aggregation thereby avoiding misconfigurations.
- SMLT client switch does not have automatic aggregation enabled (LACP disabled). SMLT aggregation switches can detect that aggregation is not enabled on the SMLT client, thus no automatic link aggregation is established until the configuration is resolved.
- Single CPU failures. In the case of a CPU failure in a system with only one switch fabric, the LACP on the other switch (or switches) detects the remote failure and triggers removal of links connected to the failed system from the LAG. This process allows failure recovery for the network along a different network path.
- LACP is not supported on SMLT NNIs because it can cause traffic to be silently dropped.

SMLT and IP routing

This section describes SMLT and IP routing interactions.

Navigation

- [SMLT and the Virtual Router Redundancy Protocol](#) on page 38
- [SMLT and VRRP BackupMaster](#) on page 38

SMLT and the Virtual Router Redundancy Protocol

Using Virtual Router Redundancy Protocol (VRRP) you can have one active primary router for each IP subnet, with all other network VRRP interfaces operating in backup mode.

The VRRP has only one active routing interface enabled. Users that access switches aggregated into two SMLT switches send their shared traffic load (based on source and destination MAC or IP addresses) on all uplinks towards the SMLT aggregation switches.

The VRRP is less efficient if you use it with SMLT. All other interfaces are in backup (standby) mode. In this case, all traffic is forwarded over the IST link towards the primary VRRP switch. All traffic that arrives at the VRRP backup interface is forwarded, so there is not enough bandwidth on the IST link to carry all the aggregated riser traffic. However, an enhancement to VRRP overcomes this issue by ensuring that the IST trunk is not used in such a case for primary data forwarding.

SMLT and VRRP BackupMaster

The VRRP BackupMaster acts as an IP router for packets destined for the logical VRRP IP address. All traffic is directly routed to the destined subnetwork and not through Layer 2 switches to the VRRP master. This avoids potential limitation in the available interswitch trunk bandwidth.

To avoid potential frame duplication problems, you can only use the VRRP BackupMaster feature for SMLT on interfaces that are defined for SMLT. You cannot use VRRP BackupMaster with hubs to avoid frame duplication or on brouter or VLAN interfaces.

When using an SMLT with routing on SMLT aggregation switches, Avaya recommends that you use VRRP for default gateway redundancy. In a VRRP environment, one switch is active and the other is backup. In an SMLT environment, you can enable the VRRP BackupMaster and use an active-active concept. The VRRP BackupMaster router routes traffic that is received on the SMLT VLAN and avoid traffic flow across the interswitch trunk. This provides true load-sharing abilities.

The BackupMaster feature provides an additional benefit. VRRP normally sends a hello packet every second. When three hello packets are not received, all switches automatically revert to master mode. This results in a 3-second outage. When you use VRRP in an SMLT environment, and a link goes down, traffic is automatically forwarded to the remaining ports configured for SMLT VRRP BackupMaster. Because both switches are processing traffic, the node immediately recognizes the VRRP state change, so there is faster failure recovery (less than 1 second).

Follow these guidelines when you use VRRP BackupMaster with SMLT:

- The VRRP virtual IP address and the VLAN IP address cannot be the same.
- Configure the hold-down timer for VRRP to a value that is approximately 150 percent of the Interior Gateway Protocol (IGP) convergence time to allow the IGP enough time to reconverge following a failure. For example, if OSPF takes 40 seconds to reconverge, configure the hold-down timer to 60 seconds.
- Stagger the hold-down timers with ARP requests. So the Avaya Ethernet Routing Switch 8800/8600 does not run ARP and SMLT at the same time causing excess CPU load. For example, if one node has the hold-down timer configured for 60 seconds, you can configure the other hold-down timer to 65 seconds.
- Enable hold-down times on both VRRP sides (Master and BackupMaster).

SMLT and SLPP

Simple Loop Prevention Protocol (SLPP) is used to prevent loops in a SMLT network. SLPP is focused on SMLT networks but works with other configurations. Avaya recommends that you always use SLPP in any SMLT environment.



Note:

When using SPBM, SLPP is not required on the B-VLAN and is not supported. The B-VLAN active topology is controlled by IS-IS that has loop mitigation and prevention capabilities built into the protocol. However, SLPP is still supported on C-VLANs to detect loops created by customers in their access networks.

SLPP requires the use of 4.1.8.2 code or higher.

Do not enable SLPP Rx on IST MLT ports. You can enable SLPP on other non-SMLT ports of any IST Core switch pair.

For information about SLPP fundamentals and configuring SLPP, see *Avaya Ethernet Routing Switch 8800/8600 Administration*, NN46205-605.

For information about SMLT and SLPP configuration examples, see the *Switch Clustering Split Multi-Link Trunking Technical Configuration Guide*, NN48500-518.

SMLT/RSMLT operational improvements

The SMLT re-architecture is supported in releases 4.1.8.2, 4.1.8.3, 5.0.1.0+, 5.1.0.0+, and 7.0.0.0+.

For previous SMLT operation, bringing the SMLTs Up/Down triggered the flushing of the entire MAC/ARP belonging to the SMLTs in both the IST core peer switches. Flushing of the MAC addresses then causes the dependent ARP to be re-resolved. For ARP resolution, ERS 8800/8600 re-ARPs for all the SMLT learned ARPs. This created a major MAC/ARP re-learning effort.

As the records are flushed, during the re-learning period the exception (learning) packets will also be continuously forwarded to the CPU, thereby increasing the CPU load. This will further slow-down the SMLT re-convergence as well as the hardware record re-programming.

Since the traffic flow is completely dependent on the hardware records, this prior behavior could adversely affect convergence times, especially in very large networks (8000+ MACs/ARPs), and those also running with many multicast streams, as multicast streams often need to be forwarded to CPU for learning, thereby also increasing CPU load. The SMLT changes improve this operation significantly, and continue to allow all previous SMLT/RSMLT topologies to be supported.

SMLT Operational Improvements will affect SMLT/RSMLT behavior in that the actual SMLT/RSMLT connection links on a powered-up IST Core Switch, will take longer to become active (link status up and forwarding) than with previous versions of code. During this time period the other Peer IST Core Switch will always be continuing to forward, therefore avoiding any loss of traffic in the network. The SMLT/RSMLT associated links will not become active upon a boot-up until the IST is completely up, and a 'sync' has occurred between the two Core IST Peer Switches.

Users may see occasional instances where the Remote SMLT Flag is False on both Peer Switches. This is normal, if the flag clears and is then set properly (False on one side, True on the other), once the FDB age-out for that associated VLAN has occurred.

For proper network "behavior" Avaya recommends to operate both IST switches with either the "new" or "old" SMLT architecture. Therefore SMLT operation between IST Peer Core switches with one switch operating with pre-4.1.8.x code, and the other operating with 4.1.8.x or later code is NOT supported.

Additionally users will see some new informational log messages generated around this behavior. These messages formats will be seen as below, under the various situations noted.

Table 1: SMLT informational message examples

Case	Message
When sync fails due to difference in IST Peer software version (pre-4.1.8.x and 4.1.8.x) where one peer supports MAC/ARP sync but	CPU5 [05/15/08 05:28:51] MLT INFO SMLT MAC/ARP Sync Requested: CAUTION do not take

Case	Message
the other does not, or some other potential issue, or maybe a mis-configuration. If the system that is reset is requesting sync, it will keep all the ports locked down until the IST comes up properly and sync has occurred. We display the error message after 5 minutes.	ANY action with the peer at this time CPU5 [05/15/08 05:33:55] MLT ERROR SMLT initial table sync is delayed or failed. Please check the peer switch for any partial config errors. All the ports in the switch except IST will remain locked

Additional SMLT considerations

Avaya does not support an additional redundant IST_MLT group between two IST peers. Use of a second MLT connection (non-IST) is supported, but such designs are not recommended.

For an IPVPN-lite configuration, where an edge 8800/8600 Cluster is configured to use a SMLT configuration toward the core (most likely square or full-mesh RSMLT), SMLT fast failover cannot always be guaranteed for this portion of the network. (CR Q2016922)

Note:

All tested failover times for SMLT/RSMLT operation are sub-second. With this release, all tested fail-back or recovery times have been improved, especially for very large scaled environments, to be within 3 seconds. This worst-case value is used in order to provide required recovery time for converged networks, especially IP phone operation. These values are for unicast traffic only; this includes voice traffic. Not all IP Multicast failover or fail-back/recovery situations can provide such times today, as many situations depend on the IPMC protocol itself recovering. For best IPMC recovery in SMLT/RSMLT designs, the use of static RPs for PIM-SM is recommended, with the same CLIP address assigned to both Core IST Peers within the Cluster, and to all switches within a full-mesh or square configuration. Failover or fail-back/recovery times for any situations that involve high-layer protocols, such as OSPF, BGP, etc., can not always be guaranteed. Reference the Network Design Guide for your specific code release for recommendations on best practices, in order to achieve best results. The 3 second recovery time is supported in Avaya SMLT/RSMLT designs which utilize VLACP. Designs that do not use VLACP or consist of a mixed vendor environment cannot be guaranteed to have a 3 second fail-back or recovery time. Avaya continues to support SMLT designs with non-Avaya devices, but recovery times for these designs cannot always be expected to be within 3 seconds.

Note:

VLACP configuration should now use values of 500 msec short timer (or higher) and a minimum timeout-scale of 5. Lower values can be used, but should any VLACP 'flapping' occur, the user will need to increase one or more of the values. These timers have been proven to work for any large scaled environments (12,000 MACs), and also provide the 3 second recovery time required for converged networks ($5 \times 500 = 2.5$ seconds). Using these values may not increase re-convergence or fail-back/recovery times, but instead guarantee these times under all extreme conditions. (CR Q01925738-01 and Q01928607) As well,

users should note that if VLACP is admin disabled on one side of the link/connection, this will cause VLACP to bring the associated remote connection down, but since the remote connection will keep link up, the side with VLACP admin disabled, will now have a black-hole connection to the remote switch, which will cause a drop of all packets being sent to it. If VLACP is disabled on one side of a connection, it MUST also be disabled on remote side or else traffic loss will likely occur. The same applies to LACP configurations for 1 port MLTs.

 Note:

LACP short timers for RSMLT/SMLT configurations must be set to 550 milliseconds or higher.

 Note:

If using VRRP with SMLT, users MUST use unique VRIDs, especially when scaling VRRP (more than 40 instances). Use of a single VRID for all instances is supported within the standard, but when such a configuration is used in scaled SMLT designs, instability could be seen. An alternative method, which allows scaling to maximum number of IP VLANs, is to use RSMLT designs instead.

 Note:

For any SMLT design, for either RSMLT L2 Edge VLANs or L2 SMLT VLANs using VRRP, it is now HIGHLY recommended to change the default VLAN FDB aging timer from its default value of 300 seconds, to now be 1 second higher than the system setting for the ARP aging timer. FDB timers are set on a per VLAN basis. If using the default system ARP aging time, config ip arp aging <minutes>, of 360 (minutes) than the proper value for the FDB aging timer, config vlan x fdb-entry aging-time <seconds>, should be 21601 seconds, which is 360 minutes (6 hours) plus 1 second. This will have the system only use the ARP aging timer for aging, versus the FDB aging timer. This value has been shown to work very well to assure no improper SMLT learning. The use of this timer has one potential side-affect: this limits the system to around a maximum of 64,000 concurrent MACs, even with timer setting. With this timer, should an edge device move, the system will still immediately re-learn and re-populate the FDB table properly, and not have to wait for the 6 hour (plus 1 second) timer to expire. No negative operational affects are known when using this timer value. For non-SMLT based VLANs the default FDB aging timer of 300 maybe used or can be changed or even also set to 21601. For this reason the default value of the FDB aging timer will remain at 300 (seconds), within all code releases.

 Note:

There are now new and improved informational and warning log messages associated with SMLT. Some messages will be generated for normal SMLT failovers, but additionally these messages can assist with troubleshooting any abnormal SMLT activity. Under abnormal conditions, please contact Avaya Technical Support for further assistance, and provide the exact log messages generated.

Virtual Link Aggregation Control Protocol

Virtual Link Aggregation Control Protocol (VLACP) is an extension to LACP used for end-to-end failure detection. VLACP is not a link aggregation protocol. It is a mechanism to periodically check the end-to-end health of a point-to-point connection. VLACP uses the Hello mechanism of LACP to periodically send Hello packets to ensure an end-to-end communication. When Hello packets are not received, VLACP transitions to a failure state, which indicates a service provider failure and that the port is disabled.

The VLACP only works for port-to-port communications where there is a guarantee for a logical port-to-port match through the service provider. VLACP does not work for port-to-multiport communications where there is no guarantee for a point-to-point match through the service provider. You can configure VLACP on a port.

VLACP can also be used with MLT to complement its capabilities and provide quick failure detection. VLACP is recommended for all SMLT access links when the links are configured as MLT to ensure both end devices are able to communicate. By using VLACP over Single-Port SMLT, enhanced failure detection is extended beyond the limits of the number of SMLT or LACP instances that can be created on an Avaya switch.

VLACP trap messages are sent to the management stations if the VLACP state changes. If the failure is local, the only traps that are generated are port linkdown or port linkup.

The Ethernet cannot detect end-to-end failures. Functions such as remote fault indication or far-end fault indication extend the Ethernet to detect remote link failures. A major limitation of these functions is that they terminate at the next Ethernet hop. They cannot determine failures on an end-to-end basis.

For example, in [Figure 10: Problem description \(1 of 2\)](#) on page 44 when the Enterprise networks connect the aggregated Ethernet trunk groups through a service provider network connection (for example, through a VPN), far-end failures cannot be signaled with Ethernet-based functions that operate end-to-end through the service provider network. The MultiLink trunk (between Enterprise switches S1 and S2) extends through the Service Provider (SP) network.

[Figure 10: Problem description \(1 of 2\)](#) on page 44 shows a MLT running with VLACP. VLACP can operate end-to-end, but can be used in a point-to-point link.




Figure 10: Problem description (1 of 2)

In the following figure, if the L2 link on S1 (S1/L2) fails, the link-down failure is not propagated over the SP network to S2 and S2 continues to send traffic over the failed S2/ L2 link.




Figure 11: Problem description (2 of 2)

Use VLACP to detect far-end failures, which allows MLT to failover when end-to-end connectivity is not guaranteed for links in an aggregation group. VLACP prevents the failure scenario.

When used in conjunction with SMLT, VLACP allows you to switch traffic around entire network devices before Layer 3 protocols detect a network failure, thus minimizing network outages.

VLACP timers and sub-100-ms core convergence

Avaya Ethernet Routing Switch 8800/8600 software uses the following VLACP timers:

- fast-periodic timer—30 to 20 000 ms; default 200 ms
- slow-periodic timer—10 000 to 30 000 ms; default 30 000 ms

Ethernet Routing Switch 8800/8600 software can attain sub-100-ms failover time. The VLACP fast-periodic timer is reduced to a minimum value of 30 ms to enable sub-100-ms convergence (period of timeout x fast-periodic timer = 3 x 30 ms = 90 ms). Sub-100-ms convergence guarantees ultra-fast convergence for critical business and multimedia applications.

To attain sub-100-ms core convergence, each switch must use the Enterprise enhanced CPU daughter card 8692SF (SuperMezz). This feature is only supported between core Ethernet Routing Switch 8800/8600 devices. For more information about SuperMezz, see *Avaya Ethernet Routing Switch 8800/8600 Installation — Modules, NN46205-304*.

Although Ethernet Routing Switch 8800/8600 ports can be configured with a VLACP fast-periodic timer of less than 30 ms, this configuration is not stable in scaled networks, and therefore not supported. Avaya recommends using VLACP with a fast-periodic timer of 30 ms and a timeout of 3 as the fastest timers in order to achieve sub-100-ms failover. For this configuration, as well as for any VLACP fast-periodic timer value of 200 ms or less, the addition of a SuperMezz is required. In addition, Avaya only supports the same minimum fast-periodic timer configuration of 30 and 3 when operating with High Availability mode enabled, no matter how scaled the environment may be.

The following figure depicts four core Ethernet Routing Switch 8800/8600 equipped with SuperMezz modules exchanging fast-periodic timer messages to achieve sub-100-ms convergence.




Figure 12: Sub-100-ms convergence between SuperMezz modules

VLACP considerations

Use the information in this section to understand the considerations when configuring VLACP into your network.

- Avaya recommends using VLACP with long timers on MLT IST links.



Note:

The switch can send VLACP PDUs at the short timeout interval, even though the long timeout interval is configured. This operation has no negative affect on VLACP operation, to other ERS 8800/8600s, or to any other Avaya ERS switches running VLACP.

- Do not use VLACP on configured LACP MLTs because LACP provides the same functionality as VLACP for link failure. VLACP and LACP running on the same link is not recommended.
- When you configure VLACP on an MLT, make sure that all MLT links on both sides are configured with the same recommended VLACP timers, and verify that VLACP is in the up state.
- In a pure Layer 2 point-to-point configuration using VLACP fast timers, 100 ms convergence is not guaranteed when ECMP is implemented on top of the link. Keep the

configuration as simple as possible with routed VLANs, but without ECMP. For fast convergence, Avaya recommends that you use RSMMLT rather than ECMP.

- If the value of VLACP timer is not set to a multiple of 10, VLACP may not function properly because the VLACP port may not be able to send out any VLACP PDUs. VLACP timers must always be set to some multiple of 10.
- When VLACP is configured, the log message LACP WARNING Received MAC-mismatched PDUs on port <port_number>, MAC <mac_address>, Please check your VLACP configuration. is sometimes observed. In the event of a VLACP misconfiguration, this message will be logged continuously in the log file, and the configuration should be corrected. If VLACP is not misconfigured and the message is observed sporadically, it should be disregarded as there is no impact to traffic.
- When a VLACP-enabled port does not receive a VLACPDU, it should enter the disabled state. There are occasions when a VLACP-enabled port does not receive a VLACPDU but remains in the forwarding state. To avoid this situation, ensure that the VLACP configuration at the port level is consistent – both sides of the point-to-point connection should be either enabled or disabled.
- Traffic loss has been observed for 1-2 seconds after a port link up occurs. This is caused by the edge switch transmitting traffic immediately upon connecting to the aggregation (IST core) switch. As soon as the link is enabled, it takes a while before VLACP engages and stops the traffic from flowing on that link.

VLACP is configured on a per port basis. The port can be either an individual port or a MLT member. VLACPDUs are sent periodically on each port where VLACP is enabled. This allows the exchange of VLACPDUs from an end-to-end perspective. If VLACPDUs are not received on a particular link, that link will be taken down after the expiry timeout occurs (timeout scale x periodic time). This implies that unless VLACP is enabled on the IST peer, the ports will stay in a disabled state. When VLACP is enabled at the IST peer, the VLACPDUs are received and the ports will be re-enabled. This behavior can be replicated despite the IST connectivity between the end-to-end peers. When you enable VLACP on the IST ports at one end of the IST, the ports are taken down along with the IST. However, the IST at the other end will stay active until the expiry timeout occurs on the other end. As soon as you enable VLACP at the other end, the VLACPDUs are received by the peer and the ports are brought up at the software level.

VLACP flap detect and damping

Link instability or packet loss can cause the Virtual Link Aggregation Control Protocol (VLACP) state of a link to toggle (flap) rapidly, bringing services (such as IP multicast) up and down in rapid succession. This behavior can cause system-wide instability, including high CPU utilization. VLACP flap detect and damping is used to automatically shut down selected VLACP links until a network administrator is able to resolve the root cause of the VLACP flapping. VLACP flap detect and damping does not support auto-recovery, therefore a network administrator must re-enable the interface manually.

When enabled, VLACP flap detect and damping shuts off a selected VLACP link if the interface flaps a specified number of times within a user-defined time frame. For example, VLACP flap

detect and damping is configured, by default, to detect 3 events within a 60 second time frame. On detection of the first event, the VLACP flap timer is started and counts off how many VLACP events occur within 60 seconds. If the number of events reaches 3 before the end of the timer, the flapping interface is shut down, the timer stops and returns to 0, and the system generates a trap and log.

 Note:

VLACP flap detect and damping is disabled by default, and should only be enabled after consultation with Avaya Client Services. Although there is some interaction with existing LACP link flap functionality, Avaya recommends that you do not use the VLACP flap dampening feature on an LACP-enabled interface.

To configure VLACP flap detect and damping using the EDM, see [Configuring VLACP flap detect and damping](#) on page 83.

To configure VLACP flap detect and damping using the CLI, see [Configuring VLACP flap detect and damping](#) on page 116.

To configure VLACP flap detect and damping using the ACLI, see [Configuring VLACP flap detect and damping](#) on page 149.

LACP and Minimum Link

The Minimum Link parameter defines the minimum number of active links required for a LAG to remain in the forwarding state. Use the Minimum-Link (MinLink) feature so that when the number of active links in a LAG is less than the MinLink parameter, the entire LAG is declared down. Prior to MinLink support, a LAG was always declared up if one physical link of the LAG was up.

Configure MinLink for each LAG; each LAG can have a different value, if required. The number of minimum links configured for an end of a LAG is independent of the other end; you can configure a different value for each end of a LAG. The default MinLink value is 1, with a range of 1 to 8.

MinLink requirements

To use MinLink on the Ethernet Routing Switch 8800/8600, you require the following:

- 8692 SF/CPU with Super Mezz or 8895 SF/CPU modules
- Software Release 4.1.5.9, or 4.1.8.0 or later

MinLink operations

If the number of active links in the LAG becomes less than the MinLink setting, the Ethernet Routing Switch 8800/8600 marks the LAG as down, and informs the remote end of the LAG state by using a Link Aggregation Control Protocol Data Unit (LACPDU). The switch continues to send LACPDUs to neighbors on each available link based on the configured timers. When the number of active links in the LAG is greater than or equal to the MinLink parameter, LACP informs the remote end, and the LAG transitions to the forwarding (up) state.

The maximum number of active links in a LAG is 8; however, it is possible to configure up to 16 links in a LAG. The eight inactive links are in Standby mode. If a link goes down, Standby links take precedence over MinLink. When an active link is disabled, the standby link with the lowest port number immediately becomes active. MinLink operates after the Standby processes finish.

Layer 3 protocols react to LAG state changes. When a LAG changes its state (goes up or down), the Ethernet Routing Switch 8800/8600 traps and logs an event. A Simple Network Management Protocol (SNMP) trap identifies that the LAG is down due to local action (local down), that the LAG is down due to remote end notification (remote down), or that the LAG is up (up).

MinLink considerations

On standard MLT links, you must enable LACP to enable MinLink.

You cannot enable MinLink on Split MultiLink Trunking (SMLT) links because the minimum number of links with SMLT can only be set to 1.

Link aggregation configuration considerations

Use the information in this section to understand the considerations and guidelines when configuring link aggregation into your network.

Navigation

- [MLT with LACP configuration considerations](#) on page 50
- [MLT with LACP and SMLT configuration considerations](#) on page 51
- [MLT with LACP and Spanning Tree configuration considerations](#) on page 53

- [LACP parameters configuration considerations](#) on page 54
- [Link aggregation and scaling](#) on page 55

MLT with LACP configuration considerations

When you configure MLT with LACP, you must enable the aggregation parameter. After you enable the aggregation parameter, the LACP aggregator maps one-to-one to the specified MultiLink trunk.

The following lists the steps that are involved when you configure MLT with LACP:

1. Assign a numeric key to the ports you want to include in the LAG.
2. Configure the LAG for aggregation.
3. Enable LACP on the port.
4. Create an MultiLink trunk and assign the same key as in step 1 to it.

The MultiLink trunk with LACP only aggregates ports whose key matches its own.

The newly created MultiLink trunk with LACP adopts the VLAN membership of its member ports when the first port is attached to the aggregator associated with this LAG. When a port detaches from an aggregator, the associated LAG port deletes the member from its list.

Prior to Release 7.1.3, you could not add or remove a VLAN from an active LACP-enabled MLT without first disabling LACP on the MLT. Starting with Release 7.1.3, this restriction is removed, namely a VLAN can be added (adding all port members) to or removed (removing all port members) from an LACP-enabled MLT. The CLI, ACLI, and EDM commands to configure this feature are shown below.

In the CLI, use either of the following commands to add a VLAN to an active LACP-enabled MLT:

- **config mlt <mlt-id> add vlan <vid>**
- **config vlan <vid> add-mlt <mlt-id>**

In the CLI, use the following command to remove a VLAN from an active LACP-enabled MLT:

- **config mlt <mlt-id> remove vlan <vid>**

In the ACLI Global Configuration Mode, use either of the following commands to add a VLAN to an active LACP-enabled MLT:

- **mlt <mlt-id> vlan <vid>**
- **vlan mlt <vid> <mlt-id>**

In the ACLI Global Configuration Mode, use the following command to remove a VLAN from an active LACP-enabled MLT:

- **no mlt <mlt-id> vlan <vid>**

In EDM, navigate to the following tab to dynamically add or remove VLANs from an active LACP-enabled MLT:

- **VLAN > MLT/LACP > Multilink/LACP Trunks**

 Note:

It is not allowed to add or remove individual ports from a VLAN when these ports are members of an active LACP-enabled MLT unless LACP is disabled first. After disabling LACP, you can use the following commands to add or remove individual ports on the MLT.

In the CLI:

- **config vlan <vid> ports add <ports in LACP-enabled MLT>**
- **config vlan <vid> ports remove <ports in LACP-enabled MLT>**

In the ACLI Global Configuration Mode:

- **vlan members <vid> <slot/ports in LACP-enabled MLT>**
- **no vlan members <vid> <slot/ports in LACP-enabled MLT>**

In EDM:

- **VLAN > Basic**

To enable tagging on ports belonging to a LAG, disable LACP on the port and then enable tagging and LACP on the port.

If you enable Open Shortest Path First (OSPF) routing on a port, do not set the LACP periodic transmission timer to less than 1 second.

MLT with LACP and SMLT configuration considerations

Split MultiLinkTrunks (SMLT) can be configured with MLT or MLT with LACP. Follow these guidelines when you configure SMLT with LACP:

- When you set the LACP system ID for SMLT, configure the same LACP SMLT system ID on both aggregation switches to avoid the loss of data. Avaya recommends that you configure the LACP SMLT system ID to be the base MAC address of one of the aggregate switches, and that you include the SMLT-ID. Ensure that the same System ID is configured on both of the SMLT core aggregation switches.
- If you use LACP in an SMLT square configuration, the LACP ports must have the same keys for that SMLT LAG; otherwise, the aggregation can fail if a switch fails.
- If an SMLT aggregation switch has LACP enabled on some of its MultiLink trunks, do not change the LACP system priority. If some ports do not enter the desired MultiLink trunk

after a dynamic configuration change, enter the following CLI command: **conf mlt <mlt-id> lacp clear-link-aggrgate**

- When you configure SMLT links, Avaya recommends that you set the multicast packets-per-second value to 6000 pps.
- Avaya recommends that you do not enable LACP on interswitch trunks to avoid unnecessary processing. Use VLACP when an optical network between the SMLT core switches requires a failure detection mechanism.
- Enabling mcast-mlt distribution on a square-mesh SMLT can cause loss of streams going through the diagonal links. This results in traffic not being distributed properly. Avaya recommends that you disable the diagonal links.

Using the SMLT system ID enables you to use any third-party switch as a wiring closet switch in an SMLT configuration. This enhancement provides an option for the administrator to configure the SMLT Core Aggregation Switches to always use the system ID. In this way, the SMLT Core Aggregation Switch will always use the same LACP key regardless of the state of SMLT Core Aggregation Switch neighbor (or the IST link). Therefore no change in LAGs should occur on the attached device. This is the case regardless of whether the device is a server or a third-party switch. This situation does not affect Avaya edge switches used in SMLT configurations. The actor system priority of LACP_DEFAULT_SYS_PRIO, the actor system ID the user configures, and an actor key equal to the SMLT-ID or SLT-ID are sent to the wiring closet switch. Avaya recommends that you configure the system ID to be the base MAC address of one of the aggregate switches along with its SMLT-id. The administrator must ensure that the same value for system ID is configured on both of the SMLT Core Aggregation Switches

The Avaya Ethernet Routing Switch 8800/8600 software does not support Simple Loop Prevention Protocol (SLPP) in an LACP–SMLT environment.

With Release 4.1.8.2 and higher, an administrator can configure the LACP SMLT System ID used by SMLT core aggregation switches. When you set the LACP system ID for SMLT, configure the same LACP SMLT system ID on both aggregation switches to avoid the loss of data.

An explanation of the importance of configuring the System ID is as follows. The LACP System ID is the base MAC address of the switch, which is carried in Link Aggregation Control Protocol Data Units (LACPDU). When two links interconnect two switches that run LACP, each switch knows that both links connect to the same remote device because the LACPDUs originate from the same System ID. If the links are enabled for aggregation using the same key, then LACP can dynamically aggregate them into a LAG (MLT).

When SMLT is used between the two switches, they act as one logical switch. Both aggregation switches must use the same LACP System ID over the SMLT links so that the edge switch sees one logical LACP peer, and can aggregate uplinks towards the SMLT aggregation switches. This process automatically occurs over the IST connection, where the base MAC address of one of the SMLT aggregation switches is chosen and used by both SMLT aggregation switches.

However, if the switch that owns that Base MAC address reboots, the IST goes down, and the other switch reverts to using its own Base MAC address as the LACP System ID. This action

causes all edge switches that run LACP to think that their links are connected to a different switch. The edge switches stop forwarding traffic on their remaining uplinks until the aggregation can reform (which can take several seconds). Additionally, when the rebooted switch comes back on line, the same actions occur, thus disrupting traffic twice.

The solution to this problem is to statically configure the same SMLT System ID MAC address on both aggregation switches.

When you install an RS or 8800 module and have SMLT and LACP configured, LACP goes down and the switch shows the message: `swA:6# CPU6 [01/09/0814:39:28] LACP INFO lacpOperDisablePort: LACP operationally disabled on port x/y because the port's capability doesn't match key z's capability.`

The solution to this problem is to globally disable and enable LACP to bring the ports back to the MLT. For reliable LACP operation, use CANA. If CANA is not supported, enforce a speed and duplex mode on all 10/100/1000 Mbps and 10/100 Mbps ports.

MLT designated ports show as null after toggling the SMLT/IST ports. This is purely a display difference. There is no impact on functionality.

MLT with LACP and Spanning Tree configuration considerations

Only the physical link state or its LACP peer status affects LACP module operation. When a link is enabled or disabled, an LACP module is notified. STP forwarding state does not affect LACP module operation. LACPDUs can be sent if the port is in an STP blocking state.

Unlike legacy MultiLink trunks, configuration changes (such as speed and duplex mode) to a LAG member port are not applied to all member ports in the MultiLink trunks. The changed port is removed from the LAG and the corresponding aggregator and the user is alerted when the configuration is created.

In contrast to MLT, MLT with LACP does not expect BPDUs to be replicated over all ports in the trunk group. Therefore, you must enter the **ntstg disable** command to disable the parameter on the Spanning Tree Group (STG) for LACP-based link aggregation. For more information about this command, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — VLANs and Spanning Tree*, NN46205-517.

ntstg disable parameter is applicable to all trunk groups that are members of the STG. This is applicable when internetworking with devices only send BPDUs out of one port of the LAG.



Important:

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

LACP parameters configuration considerations

You can configure priorities, keys, modes, and timers for the LACP.

Navigation

- [LACP priority](#) on page 54
- [LACP keys](#) on page 54
- [LACP timers](#) on page 54
- [LACP modes](#) on page 55

LACP priority

You can configure LACP priority at the system and port level as follows:

- Port priority—determines which ports are aggregated into the LAG that has more than eight ports configured to it, as in a standby-port configuration.
- System priority—generates the switch ID when communicating with other switches. For SMLT applications, use a system priority to determine a master–slave relationship between the SMLT switches. Avaya recommends that you use the default value. If you need to change it, first disable the LACP and then enable it again after you change the value.

LACP keys

LACP keys are used to determine which ports are eligible for link aggregation. The LACP keys are defined by the ports when the MultiLink trunk is configured. The ports key which match the MLT key can be aggregated into that MultiLink trunk.

- Keys do not have to match between two LACP peers.
- Keys do not have to match on SMLT core switches when you use LACP with SMLT.

LACP timers

You can customize failover times by changing the LACP timer attributes (fast periodic time, slow periodic time, and aggregate wait time). Values are set by default to match the IEEE 802.3ad values. If you change the values, they must match on the ports participating in aggregation between two devices.

Changes to LACP timer values at the global level are reflected on all ports. However, you can change the LACP timer values for each port level. When you change an LACP timer globally,

this value is set on all ports. The global timer value overwrites the local port value irrespective of the LACP state. You must configure any port values that differ from the global values.

Avaya Ethernet Routing Switch 8800/8600 software uses the following LACP timers:

- fast-periodic timer—200 to 20000 milliseconds (ms); default 1000 ms
- slow-periodic timer—10000 to 30000 ms; default 30000 ms
- aggregation-wait timer—200 to 2000; default 2000

You cannot aggregate a link if it does not receive an LACPDU for a period of timeout x slow periodic time = 3 x 30 seconds = 90 seconds. If you use the fast periodic time, the timeout period is 3 x 1000 ms = 3 seconds. You must make timer changes to all ports participating in link aggregation and to the ports on the partnering node.

Configuration changes to the LACP timers are not reflected immediately. LACP timers do not reset until the next time you restart LACP globally or on a port. This ensures consistency with peer switches.

When you enable LACP on a port, the timer values are set at the port level. You must toggle the LACP status when timer values change. This does not impact existing ports unless you toggle the LACP status on the port.

LACP modes

LACP uses two active and passive modes.

- Active mode—ports initiate the aggregation process. Active mode ports aggregate with other active mode ports or passive mode ports.
- Passive mode—ports participate in LACP but do not initiate the aggregation process. Passive mode ports must be partnered with active mode ports for aggregation to occur.

Link aggregation and scaling

The following table summarizes the number of aggregation groups supported on Ethernet Routing Switch 8800/8600 modules. Refer to the release notes included with your switch to obtain the latest scalability information.

Table 2: Link aggregation scaling for the Avaya Ethernet Routing Switch 8800/8600

Feature	Number
Aggregation groups	128
• 802.3ad aggregation groups (LAG)	
• MultiLink trunk groups	

Feature	Number
Maximum number of links (ports) for each MLT group	8

Chapter 4: Link aggregation configuration using Enterprise Device Manager

MultiLink Trunking (MLT) with Link Aggregation Control Protocol (LACP) manages switch ports and port memberships to form a link aggregation group (LAG). Configure LACP to allow dynamic bundling of physical ports to form a single logical channel.



Important:

The current release does not support standby mode for aggregation groups larger than eight ports.

Configuring global LACP parameters

Use LACP parameters to manage switch ports and their port memberships to form link aggregation groups (LAG). LACP can dynamically add or remove LAG ports, depending on their availability and states.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Double-click **MLT/LACP**
3. To enable LACP globally, select the **Enable** check box.
4. Configure the remaining parameters as desired or retain the default values.



Important:

Configuration changes to the LACP timers are not effective immediately. LACP timers are not reset until the next time LACP is restarted globally or on a port. This ensures consistency with peer switches.

5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the LACP Global dialog box.

Variable	Value
Enable	Globally enable or disable LACP.
SystemPriority	Sets the system priority for all LACP enabled aggregators and ports. The default value is 32768.
FastPeriodicTime	Indicates the number of milliseconds between periodic transmissions that use short timeouts. Sets this value to all LACP enabled ports. The range is 200–20000. The default value is 1000.
FastPeriodicTimeOper	The operating value of the fast periodic timer on the port. The default value is 1000.
SlowPeriodicTime	Indicates the number of milliseconds between periodic transmissions that use long timeouts. All LACP enabled ports get the same value from this setting. The range is 10000–30000. The default value is 30000.
SlowPeriodicTimeOper	The operating value of the slow periodic timer on the port. The default value is 30000.
AggrWaitTime	Indicates the number of milliseconds to delay aggregation to allow multiple links to aggregate simultaneously. The range is 200–2000. The default value is 2000.
AggrWaitTimeOper	The operating value of the aggregate wait timer on the port. The default value is 2000.
TimeoutScale	Assigns the value used to calculate timeout time from the periodic time. All LACP enabled ports get the same value from this setting. The range is 2–10. The default value is 3.
TimeoutScaleOper	The operating value of the timeout scale on the port. The default value is 3.
BaseMacAddr	Specifies the base MAC address of the chassis.

Variable	Value
SmltSysId	Sets the LACP system ID for SMLT. This is an optional parameter and is only used for SMLT situations. You must configure the same LACP SMLT system ID on both aggregation switches to avoid loss of data. Avaya recommends that you configure the SmltSysId so it matches the base MAC address of one of the chassis. The default value is 00:00:00:00:00:00.

Configuring LACP parameters

Configure LACP parameters to manage LACP information.



Important:

Standby mode for aggregation groups of larger than eight ports is not supported in the current release.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**
3. Click the **LACP** tab.

You can click fields to edit them but some fields are not editable.

4. Click **Apply**.

Variable definitions

Use the data in the following table to help you configure the LACP parameters.

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 six octet read-only value carrying the individual MAC address assigned to the aggregator.

Variable	Value
ActorSystemPriority	The two octet read-write value indicating the priority value associated with the actor's system ID. The default value is 32768.
ActorSystemID	The six octet read-write MAC address value used as a unique identifier for the system that contains this aggregator. From the perspective of the link aggregation mechanisms, only a single combination of actor system ID and system priority are considered. No distinction is made between the values of these parameters for an aggregator and the ports that are associated with it. The protocol is described in terms of the operation of aggregation within a single system. However, the managed objects provided for the both the aggregator and the port allow management of these parameters. The result permit a single piece of equipment to be configured by management to contain more than one system from the point of view of the operation of link aggregation. This is useful in the configuration of equipment that has limited aggregation capability.
AggregateOrIndividual	Indicates whether the aggregator represents an aggregate (true) or an individual link (false).
ActorAdminKey	The current administrative value of the key for the aggregator. The administrative key value can differ from the operational key value. This is a 16-bit read-write value. This key needs to match the LAG key. The default value is 0.
ActorOperKey	The current operational value of the key for the aggregator. The administrative key value can differ from the operational key value. This is a 16-bit read-only value. The meaning of particular key values is of local significance.
PartnerSystemID	The six 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, the value is assigned by the local system.
PartnerSystemPriority	The two octet read-only value that indicates the priority value associated with the partner system ID. If the aggregation is manually configured, this system priority value is a value assigned by the local System.
PartnerOperKey	The current operational value of the key for the aggregator current protocol partner. This is a 16 bit read-only value. If the aggregation is manually configured, the value is assigned by the local system.

Configuring LACP on a port

Configure LACP on a port to enable LACP.

Procedure steps

1. On the Device Physical View, select a port.
2. In the navigation tree, open the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the **LACP** tab.
5. To enable LACP on the port, select the **AdminEnable** check box.
6. Configure the remaining parameters as required.
7. Click **Apply**.

Variable definitions

Use the data in the following table to help you complete the LACP dialog box for the port.

Variable	Value
AdminEnable	Sets LACP status to enabled for the port. The default value is false.
OperEnable	Indicates the operational status of LACP for the port. The default value is false.
FastPeriodicTime	Specifies the number of milliseconds between periodic transmissions using short timeouts for all LACP enabled ports. The range is 200–20000. The default value is 1000.
FastPeriodicTimeOper	The operating value of the fast periodic timer on the port. The default value is 1000.
SlowPeriodicTime	Specifies the number of milliseconds between periodic transmissions using long timeouts for all LACP enabled ports. The range is 10000–30000.

Variable	Value
	The default value is 30000.
SlowPeriodicTimeOper	The operating value of the slow periodic timer on the port. The default value is 30000.
AggrWaitTime	Specifies the number of milliseconds to delay aggregation to allow multiple links to aggregate simultaneously. The range is 200–2000. The default value is 2000.
AggrWaitTimeOper	The operating value of the aggregate wait timer on the port. The default value is 2000.
TimeoutScale	Assigns the value used to calculate timeout time from the periodic time. Set this value to all LACP enabled ports. The range is 2–10. The default value is 3.
TimeoutScaleOper	The operating value of the timeout scale on the port. The default value is 3.
ActorSystemPriority	The two octet read-write value indicating the priority value associated with the actor system ID. The range is 0–65535. The default value is 32768.
ActorSystemID	The six octet read-write MAC address value used as a unique identifier for the system that contains this aggregator. From the perspective of the link aggregation mechanisms, only a single combination of actor system ID and system priority are considered, and no distinction is made between the values of these parameters for an aggregator and the ports that are associated with it; that is, the protocol is described in terms of the operation of aggregation within a single system. However, the managed objects provided for the aggregator and the port both allow management of these parameters. The result of this is to permit a single piece of equipment to be configured by management to contain more than one system from the point of view of the operation of link aggregation. This can be of particular use in the configuration of equipment that has limited aggregation capability.
ActorAdminKey	The current administrative value of the key for the aggregator. The administrative key value

Variable	Value
	can differ from the operational key value. This is a 16-bit read-write value in the range of 0–65535. This key needs to match the LAG key.
ActorOperKey	The current operational value of the key for the aggregator. The administrative key value can differ from the operational key value. This is a 16-bit read-only value. The meaning of particular key values is of local significance. The default value is 1280.
SelectedAggID	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 there is no suitable aggregator available for it to select. This value is read-only.
AttachedAggID	The identifier value of the aggregator to which this aggregation port is currently attached. Zero indicates that the aggregation port is not currently attached to an aggregator. This value is read-only.
ActorPort	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.
ActorPortPriority	The priority value assigned to this aggregation port. This 16-bit value is read-write in the range of 0–65535. The default value is 32768.
ActorAdminState	The administrative values as transmitted by the actor in LACPDUs. The values are: • lacpActive • lacpShortTimeout • aggregation
ActorOperState	The current operational values of Actor_State as transmitted by the actor in LACPDUs. This attribute value is read-only. The values are: • lacpActive • lacpShortTimeout • aggregation The default value is lacpActive.

Variable	Value
PartnerAdminSystemPriority	The current administrative value of the port number for the protocol partner. This is a 16 bit read-write value in the range of 0–65535. The assigned value is used, along with the value of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminKey, and PartnerAdminPortPriority, to achieve manually configured aggregation. The default value is 0.
PartnerOperSystemPriority	A two octet read-only value indicating the operational value of priority associated with the partner system ID. The value of this attribute can contain the manually configured value carried in PartnerAdminSystemPriority if there is no protocol partner. The default value is 0.
PartnerAdminSystemID	A six octet read-write MAC address value that represents the administrative value of the aggregation port protocol partners system ID. The assigned value is used with the values of PartnerAdminSystemPriority, PartnerAdminKey, PartnerAdminPort, and PartnerAdminPortPriority to achieve manually configured aggregation. The default value is 00:00:00:00:00:00.
PartnerOperSystemID	A six octet read-only MAC address value that indicates representing the current value of the aggregation port's protocol partner system ID. A value of zero indicates that there is no known protocol partner. The value of this attribute can contain the manually configured value carried in PartnerAdminSystemID if there is no protocol partner. The default value is 00:00:00:00:00:00.
PartnerAdminKey	The current administrative value of the key for the protocol partner. This is a 16-bit read-write value in the range of 0–65535. The assigned value is used with the value of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminPort, and PartnerAdminPortPriority to achieve manually configured aggregation. The default value is 0.
PartnerOperKey	The current operational value of the key for the aggregator current protocol partner. This is a 16-bit read-only value. If the aggregation is

Variable	Value
	manually configured, this value is assigned by the local system. The default value is 0.
PartnerAdminPort	The current administrative value of the port number for the protocol partner. This is a 16-bit read-write value in the range of 0–65535. The assigned value is used, along with the value of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminKey, and PartnerAdminPortPriority, to achieve manually configured aggregation. The default value is 0.
PartnerOperPort	The operational port number assigned to this aggregation port by the aggregation port's protocol partner. The value of this attribute can contain the manually configured value carried in AggPortPartnerAdminPort if there is no protocol partner. This 16 bit value is read-only. The default value is 0.
PartnerAdminPortPriority	The current administrative value of the port priority for the protocol partner. This is a 16-bit read-write value in the range of 0–65535. The assigned value is used with the values of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminKey, and PartnerAdminPort to achieve manually configured aggregation. The default value is 0.
PartnerOperPortPriority	The priority value assigned to this aggregation port by the partner. The value of this attribute can contain the manually configured value carried in PartnerAdminPortPriority if there is no protocol partner. This 16 bit value is read-only. The default value is 0.
PartnerAdminState	A string of eight bits, corresponding to the current administrative value of Actor state for the protocol partner. This attribute value is read-write. The assigned value is used to achieve manually configured aggregation. The values are: • lacpActive • synchronization • defaulted

Variable	Value
	• lacpShortTimeout • collecting • expired • aggregation • distributing
PartnerOperState	This value reflects the value of PartnerAdminState. This attribute value is read-only.

Chapter 5: MLT and SMLT link aggregation configuration using Enterprise Device Manager

Configure link aggregation to provide link level redundancy and increase load sharing. Multi-Link Trunking (MLT) is a link aggregation technology that allows grouping several physical Ethernet links into one logical Ethernet link to provide fault-tolerance and high-speed links between routers, switches, and servers. SMLT is an option that improves Layer 2 (bridged) resiliency.

Adding a MultiLink/LACP trunk

Add a MultiLink/LACP trunk to configure a trunk.



Important:

Ensure that all ports that belong to the same MLT/LACP group use the same port speed, for example, 1 Gbit/s, even if auto-negotiation is used. This is a requirement that is not enforced by the software.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Double-click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. Click **Insert**.
5. In the Id box, type the ID number of the MLT.
6. In the **PortType** section, select **trunk**.
7. In the Name box, type a name for the MLT, or accept the default.
8. In the PortMembers box, click [...].
9. Select the desired ports in the Port editor: Port Members dialog box.
10. Click **Ok**
11. In the VlanIds box, click [...].

12. Select the desired VLANs in the VlanIds dialog box.
13. Click **Ok**.
14. In the **MltType** section, select the MLT type.

If you select splitMLT, enter the SMLTID number in the SmltID dialog box that appears.
15. In the Multicast Distribution box, select one of the following:
 - **enable**
 - **disable**
16. Select or clear the **NtStgEnable** check box.
17. In the Aggregatable box, select one of the following:
 - **enable**
 - **disable**
18. In the AggMinLink box, configure the number of active links required for this LAG to remain in the forwarding state.
19. Click **Insert**.

The MLT is added to the MultiLink/LACP Trunks tab in the MLT_LACP box.

Variable definitions

Use the data in the following table to help you configure the MultiLink/LACP Trunks dialog box.

Variable	Value
Id	Uniquely identifies a value in the MLT.
PortType	Sets the value to access or trunk port. When the aggregatable field is set to enable, this field is read-only. The default value is access.
Name	Indicates the name given to the MLT.
PortMembers	The ports assigned to the MLT. 10Base-T, 100Base-TX, 100Base-FX, and Gigabit Ethernet ports support MLT. All ports in an MLT must have the same settings for speed and duplex, but can have different media types. All untagged ports must belong to the same STG. For Ethernet Routing Switch 8800/8600 modules, up to eight same-type ports can belong to a single MLT. When the aggregatable field is set to enable, this field becomes read-only.

Variable	Value
	 Important: Ensure that all ports that belong to the same MLT/LACP group use the same port speed, for example, 1 Gbit/s, even if auto-negotiation is enabled.
VlanIds	Indicates to which ports the VLANs belong. When the aggregation field is set to enable, this field is read-only.
MltType	Specifies the type of MLT: • normalMLT (default) • istMLT • splitMLT
RunningType	SMLT running type. The default value is normalMLT.
SmltId	The SMLT ID assigned to both ends of the split trunk. The corresponding SMLT between aggregation switches must have the same SMLT ID.
AggMinLink	Configures the minimum number of active links (MinLink) required for the LAG to remain in forwarding state. If the number of active links falls below this number, the Ethernet Routing Switch 8800/8600 declares the entire LAG down. The default is 1, and the range is 1 to 8.
IfIndex	Interface index.
Multicast Distribution	The multicast distribution feature state on MLT ports: • enable • disable (default) You must also configure Multicast distribution on the chassis (Edit, Chassis, Mcast MLT Distribution). For more information, see <i>Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Multicast Routing Protocols, NN46205-501</i> .
ClearLinkAggregate	Disables and re-enables aggregation on the MLT.
NtStgEnable	Specifies if this MLT is operating in Avaya Mode or in Cisco Mode. • true—Avaya mode (default) • false—Cisco mode
DesignatedPort	Indicates the designated port the switch uses to send multicast traffic to the MLT, if the multicast distribution feature is disabled.
Aggregatable	Specifies whether link aggregation is enabled or disabled.

Variable	Value
AggOperState	Indicates the link aggregation state on an MLT.
AggTimeofLastOperChange	Indicates the time value since the interface entered its current operational state.

Adding ports to an MLT

Add ports to an MLT to insert MultiLink/LACP trunks.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. Click **Insert**.
5. In the PortMembers box, click [...] to add ports.
6. Do one of the following:
 - Select the ports from the Port Editor: PortMembers dialog box to add to the PortMembers box.
 - Click **All** to add all ports to the MLT.

For Ethernet Routing Switch 8800/8600 modules, up to eight same-type ports can belong to a single MLT.
7. Click **Insert**.

The port numbers are added to the selected MLT on the MultiLink/LACP Trunks tab in the MLT/LACP box.
8. Click **Apply**.

The ports are added to the MLT.

Adding an MLT-based SMLT

Add an MLT-based SMLT.

**Important:**

For Ethernet Routing Switch 8800/8600 modules, up to eight same-type ports can belong to an MLT.

Prerequisites

- Create an MLT before creating an SMLT. You can create an SMLT by selecting the MLT type as SMLT and then specifying an SMLT ID.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. Click **Insert**.
5. In the Id box do one of the following:
 - type an available MLT ID number
 - use the default
6. In the PortType box, select the port type as one of the following:
 - **Access**
 - **Trunk**
7. In the Name box, type a name to identify the MLT-based SMLT port.
8. In the PortMembers box, click [...].
9. Click the ports you want to include in the MLT-based SMLT.
10. Click **Ok**.

The ports are added to the PortMembers box on the Insert MultiLink/LACP Trunks tab.

11. In the VlanIds box, click [...].
12. Select the VLAN IDs for the MLT-based SMLT port, and then click **Ok**.

The VLANs are added to the VlanIds box in the MLT, Insert Trunks box.
13. In the MltType box, select **splitMLT**.

The SmltId box appears.

14. In the SmltId box, type an unused SMLT ID.

The corresponding SMLT between aggregation switches must have matching SMLT IDs. The same ID number must be configured on both switches.

15. In the Multicast Distribution box, select one of the following:

- **enable**
- **disable**

16. Select or clear **NtStgEnable** check box.

17. In the Aggregatable box, select one of the following:

- **enable**
- **disable**

18. In the AggMinLink box, configure the number of active links required for this LAG to remain in the forwarding state.

19. Click **Insert**.

The new MLT-based SMLT appears in the MultiLink/LACP Trunks tab.

20. Click **Close**.

The MLT-based SMLT is added.

Viewing MLT-based SMLTs

Use MLT-based SMLTs to view the MLT-based SMLT configuration.



Note:

If you disable and enable the SMLT or IST ports, the designated port information may not appear correctly.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Click **SMLT**.
3. Click the **SMLT Info** tab.

The SMLT Info tab appears with all the configured MLT-based SMLT displayed.

Variable definitions

Use the data in the following table to help you view the MLT-based SMLT configuration.

Variable	Value
Id	Shows the MLT ID for this SMLT Read-only.
SmltId	Shows the MLT-based SMLT ID number.
MltType	Specifies the type of MLT: • normalMLT • istMLT • splitMLT
RunningType	Indicates the MLT operational type: • normalMLT • istMLT • splitMLT

Adding ports to an MLT-based SMLT

Use the MultiLink/LACP Trunks tab to add ports to the MLT-based SMLT.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. Click **Insert**.
5. In the PortMembers box, click [...] for the MLT-based SMLT to which you are adding ports.
6. Do one of the following:
 - Select the port numbers to be added.
 - Click **All** to select all ports.



Important:

For Ethernet Routing Switch 8800/8600 modules, up to eight same-type ports can belong to a single MLT.

7. Click **Ok**.

The ports are added to the Port Members box on the MLTs tab.

8. Click **Apply**.

The ports are added to the MLT-based SMLT.

Configuring an IST MLT

Configure an IST MLT to add IST to MLT_LACP.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**
3. Click the **MultiLink/LACP Trunks** tab.
4. Click **Insert**.
5. In the PortMembers box for the IST MLT, click [...].

The **Port Editor: PortMembers** dialog box appears showing the available ports.

6. Click the ports you want to include in the IST MLT.
7. Click **Ok**.

The ports are added to the PortMembers box for the IST MLT in the Insert MLTs tab.

8. In the MltType box, select **istMLT**.
 9. In the PortType box, select **trunk**.
 10. Configure the remaining boxes as required.
 11. Click **Insert**.
- The IST MLT dialog box appears.
12. Type the peer IP address and the VLAN ID in the box.
 13. Select **enable**.
 14. Click **Apply**.

The IST is added to the MLT/LACP box.

15. In the Device Physical View, select a port on the Ethernet module.
16. In the navigation tree, open the following folders: **Configuration > Edit > Port**.
17. Click **General**.
18. Click the **CP Limit** tab.
19. In the CpLimitEnable option for each port, select **disable**.
20. Click **Apply**.

Variable definitions

Use the data in the following table to help you configure the IST MLT dialog box.

Variable	Value
PeerIp	An IST MLT peer IP address.
VlanId	An IST VLAN ID number.
StpEnable	Enables and disables Stp functionality.
SessionEnable	Enables and disables IST functionality. The default value is disable.
SessionStatus	Displays the status of the IST session.

Editing an IST

Edit an IST to make changes to an IST MLT.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Double-click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. Select the IST.
5. Click **IstMlt**.
6. In the PeerIp box, enter the peer IP address.
7. In the VlanId box, enter a VLAN ID.

8. In the SessionEnable box, select one of the following:

- **enable**
- **disable**

9. Click **Apply**.

The IST MLT dialog box closes and the changes are applied.

Enabling SMLT Interoperability with MSTP using EDM

Use the following procedure to enable SMLT Interoperability for MSTP.

Prerequisites

- Configure your switch to MSTP mode
- Configure an IST MLT on your switch

Procedure steps

1. From the Configuration tree, click **VLAN**.
2. From the VLAN tree, click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. From the list, click an **IST MLT**.
5. On the toolbar, click the **IstMlt** button.
6. In the StpEnable box, click the **enable** button.
7. On the toolbar, click **Apply**.

Disabling SMLT Interoperability with MSTP using EDM

Use the following procedure to disable SMLT Interoperability for MSTP.

Procedure steps

1. From the Configuration tree, click **VLAN**.
2. From the VLAN tree, click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. From the list, click an **IST MLT**.
5. On the toolbar, click the **IstMLT** button.
6. In the StpEnable box, click the **disable** button.
7. On the toolbar, click **Apply**.

**Note:**

If you want to disable SMLT Interoperability for MSTP, ensure that you remove redundant paths, for VLANs, to the existing IST path. For example, you should not break or remove an MSTP device chain until you disable SMLT interoperability. Otherwise a bridging loop forms and your network fails.

Displaying IST MLT information using EDM

Use the following procedure to display IST MLT information.

Procedure steps

1. From the Configuration tree, click **VLAN**.
2. From the VLAN tree, click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.

Configuring a single port SMLT

Configure a split MultiLink trunk that uses a single port. Use single port SMLT to permit the scaling of the number of SMLT on a switch to the maximum number of available ports.

**Important:**

Only VLACP is supported on a single port SMLT.

Prerequisites

- If the port is configured as MLT or MLT-based SMLT, you cannot configure a single port SMLT.

Procedure steps

1. On the Device Physical View, select a port.
2. In the navigation tree, open the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the **SMLT** tab.

 Important:

The SMLT tab indicates if this port is configured as MLT or MLT-based SMLT. If so, you cannot configure a Single Port SMLT.

5. Click **Insert**.
6. In the SmltId dialog box, enter an unused SMLT ID number.
7. Click **Insert**.

A warning message appears, informing you that the STP has been disabled while configuring the port with SMLT.

8. Click **Ok**.

The SMLT ID is entered.

Variable definitions

Use the data in the following table to help you configure the SmltId dialog box.

Variable	Value
Port	The slot/port number of the port.
MltId	The MLT ID. • A value of 1–128 indicates that the port is part of an MLT, and Single Port SMLT cannot be configured on this port. • A value of 0 indicates that no MLT is assigned, and the port can be configured for Single Port SMLT.

Variable	Value
SmltId	The SMLT ID in the range 1–512. • A read-only field indicates the Single Port SMLT ID assignment for the port. • A blank field indicates the port is not configured for Single Port SMLT. Find an unused SMLT ID by showing the currently used IDs.

Viewing Single Port SMLTs

Use the Single Port SMLTs tab to view the single port SMLTs configured on the switch.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Click **SMLT**.
3. In the Single Port SMLT tab, the single port SMLTs configuration information appears.

Variable definitions

Use the data in the following table to help you view the single port SMLTs.

Variable	Value
Port	The port interface index number.
SmltId	The ID number of the single port SMLT.
RunningType	The port operational type: • normalMLT • istMLT • splitMLT

Deleting a Single Port SMLT

Use the SMLT tab to delete the single port SMLT.

Procedure steps

1. On the Device Physical View, select a port.
2. In the navigation tree, open the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the **SMLT** tab.
5. Select the SMLT that you want to delete.
6. Click **Delete**.
7. Click **Ok** in the confirm message box.

The single port SMLT is deleted.

Chapter 6: VLACP configuration using Enterprise Device Manager

Configure Virtual LACP (VLACP) to implement link status control protocol at the port level. Note that VLACP cannot interoperate with LACP.

Enabling VLACP globally

Enable VLACP globally to detect for end-to-end failure. VLACP uses the Hello mechanism of LACP to periodically send Hello packets to ensure there is an end-to-end approach. When Hello packets are not received, the VLACP transitions to a failure state and the port is disabled.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **VLACP** tab.
4. In the **VLACP** tab, select the **VlACPEnable** check box to enable VLACP globally.
5. Click **Apply**.

Configuring VLACP on a port

Enable VLACP on a port. VLACP periodically checks the end-to-end health of a point-to-point connection.

Procedure steps

1. On the Device Physical View, select a port.
2. In the navigation tree, open the following folders: **Configuration > Edit > Port**.
3. Double-click **General**.
4. Click the **VLACP** tab.
5. In **VLACP** tab, select the **AdminEnable** check box to enable VLACP on the port.
6. Configure the remaining fields as required.
7. Click **Apply**.

Variable definitions

Use the data in the following table to help you configure VLACP on a port.

Variable	Value
AdminEnable	Assigns the enabled status for VLACP for the port.
OperEnable	Indicates the operational status of VLACP for the port. The default value is false.
FastPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using short timeouts. Sets all LACP enabled ports the same value from this setting. The range is 10–20000. The default value is 200.
SlowPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using long timeouts. Sets all LACP enabled ports the same value from this setting. The range is 10000–30000. The default value is 30000.
Timeout	The timeout control value. Specify long or short timeout. The default value is long.
TimeoutScale	Assigns the value used to calculate timeout time from the periodic time for all VLACP enabled ports. $\text{Timeout} = \text{PeriodicTime} \times \text{TimeoutScale}$. The range is 2–10. The default value is 3.
EtherType	The VLACP protocol identification. The ID is in hexadecimal format. The default value is 0x8103.
EtherMacAddress	The multicast MAC address exclusively used for VLACPDUs.

Variable	Value
PortState	Specifies the VLACP port state.

Configuring VLACP flap detect and damping

Configure VLACP flap detect and damping to automatically shut down a VLACP link that is flapping.

Prerequisites

- Make sure that the VLACP flap interval is long enough to recognize that a port is flapping. (The default settings may not allow enough time.) A VLACP down event does not occur until the receiving side has not seen a VLACP PDU for a number of seconds (**Timeout** x **TimeoutScale**). Then this has to happen several times (**FlapFrequency**) within a specified period of time (**FlapInterval**).

To ensure that the interval is long enough, use the following formula:

$$\text{FlapInterval} > \text{FlapFrequency} \times \text{Timeout} \times \text{TimeoutScale}$$

The **FlapInterval** and **FlapFrequency** commands are described in this procedure.



Note:

VLACP flap detect is also triggered when there is a physical link down event on the underlying Ethernet port. This replaces the link-flap-detect global functionality, which is by default enabled.

For information on the **Timeout** and **TimeoutScale** commands, see [Configuring VLACP on a port](#) on page 81.

Procedure steps

1. On the Device Physical View, select a port.
2. In the navigation tree, open the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the **VLACP** tab.
5. Configure the remaining fields as required.
6. Click **Apply**.

Variable definitions

Use the data in the following table to help you to configure VLACP flap detect and damping.

Variable	Value
FlapDetect	Enables VLACP flap detect and damping. The default is disable.
FlapFrequency	Specifies the flap frequency. The value ranges from 3–30. The default is 3.
FlapInterval	Specifies the flap interval, in seconds. The value ranges from 10–600. The default is 60.

Chapter 7: Link aggregation configuration using the CLI

MultiLink Trunking (MLT) with Link Aggregation Control Protocol (LACP) manages switch ports and port memberships to form a link aggregation group (LAG). Configure LACP to allow dynamic bundling of physical ports to form a single logical channel.

LACP can be described in terms of link aggregation operations within a single system. You can configure a single piece of equipment so it contains more than one system (from the point of view of the link aggregation operation).

Prerequisites

- Changes to LACP made at the global level override and reset all port level settings.



Important:

When you globally set the LACP system priority, it applies to all LACP-enabled aggregators and ports.

When you enable LACP on an aggregator or port, it uses the global system priority value.

- If you make a timer change, restart LACP (globally or on the port) so the changes are consistent across the link.



Important:

Configuration changes to LACP timers are not reflected immediately.

LACP timers are not reset until the next time LACP is restarted globally or on a port. This ensures consistency with peer switches.

- Standby mode for LACP aggregation groups of larger than eight ports is not supported in this release.

Job aid: Roadmap of LACP commands

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
config lacp	aggr-wait-time <milliseconds>
	disable
	enable
	smlt-sys-id <BaseMac>
	system-priority <integer>
	fast-periodic-time <milliseconds>
	slow-periodic-time <milliseconds>
	timeout-scale <integer>
	info
config <port-type> <slot port> lacp	disable
	enable
	aggr-wait-time <milliseconds>
	fast-periodic-time <milliseconds>
	info
	key <integer>
	aggregation <true false>
	mode <active passive>
	partner-key <int>
	partner-port <int>
	partner-port-priority <int>
	partner-state <hex>
	partner-system-id <mac>
	partner-system-priority <int>
	port-priority <integer>
	slow-periodic-time <milliseconds>
	system-priority <integer>

Command	Parameter
	system-priority <integer>
	timeout {long short}
	timeout-scale <integer>
show lacp info	
show ports info lacp	actor-admin [vlan <value>] [port <value>]
	all [vlan <value>] [port <value>]
	actor-oper [vlan <value>] [port <value>]
	extension [vlan <value>] [port <value>]
	partner-admin [vlan <value>] [port <value>]
	partner-oper [vlan <value>] [port <value>]
show ports stats lacp [port <value>]	

Configuring global LACP parameters

Configure LACP parameters to enable or disable LACP globally. When the LACP system priority is set globally, it applies to all LACP-enabled aggregators and ports. When you enable LACP on an aggregator or a port, it will use the global system priority value.



Important:

Changes made at the global level override and reset all port level settings.

Procedure steps

1. Configure LACP globally by using the following command:

```
config lacp
```

Variable definitions

Use the data in the following table to help you use the **config lacp** command.

Variable	Value
aggr-wait-time <milliseconds>	Sets the aggregation wait time (in milliseconds) for a specific port type. The default value is 2000.
disable	Disables LACP for a specific port type.
enable	Enables LACP for a specific port type. The default value is disable.
fast-periodic-time <milliseconds>	Sets the fast-periodic time (in milliseconds) for a specific port type. The default value is 1000.
info	Shows the current level parameter settings and next level directories.
smlt-sys-id <baseMAC>	Sets the LACP system ID globally. Enter a MAC address in the following format: 0x00:0x00:0x00:0x00:0x00:0x00.
system-priority <integer>	Sets system priority for a specific port-type. <i>integer</i> is the system priority value. The integer value ranges from 0 to 65535. The default value is 32768.
slow-periodic-time<milliseconds>	Sets the slow periodic time globally. <i>milliseconds</i> is the slow periodic time value. The default value is 30000.
timeout-scale <integer>	Sets a timeout scale for a specific port type. The default value is 3, and the value ranges from 2 to 10.

Configuring LACP on a port

Configure LACP on a port to enable or disable LACP on the selected ports.

A port can operate in active or passive mode. You can configure a port to be an individual link or an aggregatable link.

**Important:**

Changes made at the global level override and reset all port level settings.

Procedure steps

1. Configure LACP on a port by using the following command:

```
config {ethernet|pos} <portList> lacp
```

Specify ethernet for Ethernet ports; pos for Packet over SONET ports.

Variable definitions

Use the data in the following table to help you use the **config {ethernet|pos} <portList> lacp** command.

Variable	Value
aggregation <true false>	Sets individual port or aggregation for this port. • true sets the port as aggregation . • false sets the port as individual. The default value is false.
aggr-wait-time <milliseconds>	Sets the aggregation wait time (in milliseconds) for this port. The default value is 2000.
disable	Disables LACP for this port.
enable	Enables LACP for this port. The default value is disable.
fast-periodic-time <milliseconds>	Sets the fast-periodic time (in milliseconds) for this port. The default value is 1000.
info	Shows the current level parameter settings and next level directories.
key <integer>	Sets LACP aggregation key for this port. You can use a default key only for individual ports.
mode <active passive>	Sets active or passive mode for this port. The default value is passive.
partner-key <int>	The LACP partner administrative key. The integer value ranges from 0 to 65535. The default value is 0.
partner-port <int>	Sets the port-partner administrative port value. The integer value ranges from 0 to 65535.

Variable	Value
	The default value is 0.
partner-port-priority <int>	Sets the port-partner administrative port priority value. The integer value ranges from 0 to 65535. The default value is 0.
partner-state <0x00–0xff>	Sets the partner administrative state bitmask. Specify the partner administrative state bitmap in the range 0x00–0xff. The bit to state mapping is Exp, Def, Dis, Col, Syn, Agg, Time, and Act. For example, to set the two partner-state parameters • Act = true • Agg = true • Specifies a value of 0x05 (bitmap = 00000101).
partner-system-id <mac>	Sets the partner administrative system ID. • <i>mac</i> is the LACP partner administrative system ID MAC address in the format: 0x00:0x00:0x00:0x00:0x00:0x00.
partner-system-priority <int>	Sets the partner administrative system priority value. • <i>int</i> is the LACP partner administrative system priority. The integer value ranges from 0 to 65535. The default value is 0.
port-priority <integer>	Sets the port priority. The default value is 32768. • <i>integer</i> is the port priority value. The integer value ranges from 0 to 65535.
slow-periodic-time <milliseconds>	Sets the slow periodic time (in milliseconds) for a specific port type. The default value is 30000.
system-priority <integer>	Sets system priority for this port. • <i>integer</i> is the system priority value. The integer value ranges from 0 to 65535. The default value is 32768.
timeout {long short}	Sets the timeout value to either long or short for this port. The default value is long.
timeout-scale <integer>	Sets a timeout scale for this port. The default value is 3, and the value ranges from 2 to 10.

LACP show commands

Use the LACP show commands to displays LACP configuration information.

Navigation

- [Viewing global LACP configuration information](#) on page 91
- [Viewing LACP configuration information](#) on page 91

Viewing global LACP configuration information

View global LACP configuration information so that you can display global LACP configuration information.

Procedure steps

1. View global LACP configuration information by using the following command:

```
show lacp info
```

Viewing LACP configuration information

View LACP configuration information to determine the LACP parameters and to ensure your configuration is correct.

Procedure steps

1. View LACP configuration information by using the following command:

```
show ports info lacp
```

Variable definitions

Use the data in the following table to help you use the **show ports info lacp** command.

Variable	Value
actor-admin [vlan <value>] [port <value>]	Shows port LACP actor administrative information.

Variable	Value
	• vlan <i>value</i> is the VLAN ID. • port <i>value</i> is the port or port list.
actor-oper [vlan <value>] [port <value>]	Shows port LACP actor operational information. • vlan <i>value</i> is the VLAN ID. • port <i>value</i> is the port or port list.
all [vlan <value>] [port <value>]	Shows all port LACP information. • vlan <i>value</i> is the VLAN ID. • port <i>value</i> is the port or port list.
extension [vlan <value>] [port <value>]	Shows port LACP timer information. • vlan <i>value</i> is the VLAN ID. • port <i>value</i> is the port or port list.
partner-admin [vlan <value>] [port <value>]	Shows port LACP partner administrative information. • vlan <i>value</i> is the VLAN ID. • port <i>value</i> is the port or port list.
partner-oper [vlan <value>] [port <value>]	Shows port LACP partner operational information. • vlan <i>value</i> is the VLAN ID. • port <i>value</i> is the port or port list.

Chapter 8: MLT and SMLT link aggregation configuration using the CLI

Configure link aggregation to provide link level redundancy and increase load sharing. Multi-Link Trunking (MLT) is a link aggregation technology that allows grouping several physical Ethernet links into one logical Ethernet link to provide fault-tolerance and high-speed links between routers, switches, and servers. SMLT is an option that improves Layer 2 (bridged) resiliency.

Job aid: Roadmap of MLT and SMLT commands

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
config mlt <mid>	create
	cp-limit <enable disable> [multicast-limit <value>] [broadcast-limit <value>]
	delete
	info
	mcast-distribution <enable disable>
	name <string>
	ntstg <enable disable>
	perform-tagging <enable disable>
config mlt <mid> add	info
	ports <ports>
	vlan <vid>
config mlt <mid> remove	info
	ports <ports>

Command	Parameter
	vlan <vid>
config mlt <mlt id> lacp	clear-link-aggrgate
	disable
	enable
	info
	key <integer>
	min-link <integer>
	system-priority <integer>
config mlt <mid> smlt	create smlt-id <value>
	delete
	info
config mlt <mid> ist	create ip <value> vlan-id <value>
	delete
	disable
	enable
	info
	spanning-tree <enable disable>
config ethernet <slot/port> cp-limit	broadcast-limit <value>
	<enable disable>
	multicast-limit <value>
config ethernet <port> smlt <smltid>	create
	delete
	info
config sys set smlt-on-single-cp <enable disable> [timer <value>]	
show mlt error collision [<mid>]	

Command	Parameter
show mlt error main [<mid>]	
show mlt info [<mid>]	
show mlt ist info	
show mlt lacp info <ifindex>	
show mlt show-all [file <value>]	
show mlt stats [<mid>]	
show ports info smlt [vlan <value>] [port <value>]	
show ports info config [vlan <value>] [port <value>]	
show smlt	

Configuring an MLT

Configure an MLT to set up MLTs on the switch.

Procedure steps

1. Configure an MLT by using the following command:

```
config mlt <mid>
```

Variable definitions

Use the data in the following table to help you use the **config mlt <mid>** command.

Variable	Value
create	Creates a LAG.
cp-limit {enable disable} [multicast-limit <value>] [broadcast-limit <value>]	Sets the control packet rate limit. • enable disable enables or disables control packet rate limit. The default value is enable. To reenabling the ports, issue the command config ethernet slot/

Variable	Value
	port state disable, and then enable. • multicast-limit <value> is the multicast control frame rate. The default value is 10000. • broadcast-limit <value> is the broadcast frame rate. The default value is 10000.
delete	Deletes a LAG.
info	Shows current settings for the specified LAG.
mcast-distribution {enable disable}	Enables or disables multicast distribution per LAG. Multicast distribution is disabled by default. For more information, see <i>Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Multicast Routing Protocols</i>, NN46205-501.
name <string>	Names a LAG. • <i>string</i> is the name, from 0 to 20 characters.
ntstg {enable disable}	Enables or disables NTSTG. The default value is enable.
perform-tagging {enable disable}	Enables or disables tagging on a link aggregation port. The default value is disable.

Adding ports to an MLT LAG

Add ports to an MLT LAG to add an existing VLAN to a link aggregation configuration.

Procedure steps

1. Adding ports to an MLT LAG by using the following command:

```
config mlt <mid> add
```

Variable definitions

Use the data in the following table to help you use the **config mlt <mid> add** command.

Variable	Value
info	Shows ports and VLANs added to the LAG.
ports <ports>	Adds ports to the LAG. • <ports> is the port number or a list of ports you want to add to the LAG. If the port you are configuring already has an SMLT ID on it, you cannot add it to the LAG.
vlan <vid>	Adds an existing VLAN to the LAG. • <vid> is the VLAN ID or a list of VLAN IDs you want to add to the LAG. The range is 1–4094.

Procedure job aid: Example of adding ports to an MLT LAG

To add the ports to the LAG use the following command:

```
ERS-8600:5#/config/mlt/1# add port 1/1-1/7,1/9
ERS-8600:5#/config/mlt/1# add vlan 1-8,10
ERS-8600:5#/config/mlt/1# info
Sub-Context: Current Context: ports : 1/1-1/7,1/9 vlan : 1-8,10
```

Removing ports from an MLT LAG

Remove ports from an MLT LAG to remove a VLAN from an MLT configuration.

Prerequisite

- Before removing a port member from an MLT, you need to first disable the port. This ensures that the other side will bring its corresponding port member down. This achieves parity on both sides and avoids traffic disruptions.

Procedure steps

1. Remove ports from an MLT LAG by using the following command:

```
config mlt <mid> remove
```

Variable definitions

Use the data in the following table to help you use the **config mlt <mid> remove** command.

Variable	Value
info	Shows the ports and VLANs removed from the MLT.
ports <ports>	Removes ports from the MLT. • <i>ports</i> is the port number or a list of ports you want to remove from the MLT.
vlan <vid>	Removes a VLAN from the MLT. • <i>vid</i> is the VLAN ID or a list of VLAN IDs you want to remove from the LAG. The range is 1–4094.

Viewing MLT LACP configuration information for aggregators

View the MLT LACP configuration information per aggregator to obtain LACP configuration information per aggregator.

Procedure steps

1. View MLT LACP configuration information per aggregator by using the following command:

```
show mlt lacp info <ifindex>
```

Configuring LACP on an MLT

Configure an MLT with LACP to use the dynamic link aggregation function.

Prerequisites

- To attach ports to an aggregator, their system priorities must be the same; otherwise, they are considered to be operating in two different switches.
- To attach ports to an aggregator, the actor admin key must match the key of the LAG.

Procedure steps

1. Configure LACP on an MLT by using the following command:

```
config mlt <mlt id> lacp
```

Variable definitions

Use the data in the following table to help you use the **config mlt <mlt id> lacp** command.

Variables	Value
clear-link-aggregate	Clears link aggregation for a specific MLT. Using clear-link-aggregate is equivalent to disabling and re-enabling LACP on the multilink trunk.
disable	Disables LACP for a specific MLT. When disabled, the multilink trunk functions as a legacy multilink trunk.
enable	Enables LACP for a specific MLT.

Variables	Value
info	Shows current level parameter settings and next level directories.
key <integer>	Sets LACP aggregator key for a specific MLT. • <i>integer</i> is the LACP key.  Important: If you set the key to 0, any LACP-enabled port with a non-matching key will be allowed to join the LAG.
min-link <integer>	Configures the minimum number of active links (MinLink) required for the LAG to remain in forwarding state. If the number of active links falls below this number, the Avaya Ethernet Routing Switch 8800/8600 declares the entire LAG down. The default value is 1.
<mlt id>	Specifies the MLT ID.
system-priority <integer>	Sets LACP system priority for a specific MLT. • <i>integer</i> is the system priority within the range 0–65535. The default value is 32768.

Creating a SMLT from an existing MLT

Create a SMLT from an existing MLT to split physical ports between two switches, which provides active load sharing and high availability.

Prerequisites

- Create an MLT before you create a split in the MLT.

Procedure steps

1. Create a SMLT from an existing MLT by using the following command:

```
config mlt <mid> smlt
```

Variable definitions

Use the data in the following table to help you use the **config mlt <mid> smlt** command.

Variable	Value
create smlt-id <value>	Creates a split MultiLink trunk from an existing MultiLink trunk. • <i>value</i> is an integer value with a range of 1 to 256. The value must match the peer switch SMLT-ID. If the split MultiLink trunk, ID already exists on a single port split MultiLink trunk, you cannot assign it to an MLT-based split MultiLink trunk.
delete	Deletes an existing SMLT.
info	Shows ports and VLANs added to the MultiLink trunk.
<mid>	Specifies the MLT ID.

Creating an interswitch trunk

Create an interswitch trunk from an existing MLT.

Procedure steps

1. Create an interswitch trunk to create an interswitch trunk by using the following command:

```
config mlt <mid> ist
```

Variable definitions

Use the data in the following table to help you use the **config mlt <mid> ist** command.

Variable	Value
create ip <value> vlan-id <value>	Creates an interswitch trunk from an existing MLT. • ip value is a peer IP address • vlan-id <value> is an integer value. The peer IP address is the IP address of the IST VLAN on the other aggregation switch.
delete	Deletes an existing IST. You must disable an IST before you can delete it.
disable	Disables an existing IST.
enable	Enables an existing IST.
info	Shows current level parameter settings and next level directories.
<mid>	Specifies the MLT ID.

Creating an interswitch trunk from an existing MLT

Create an interswitch trunk from an existing MLT to create an IP VLAN id.

Procedure steps

1. Create an interswitch trunk from an existing MLT by using the following command:

```
config mlt <mlt-id> ist create ip <value> vlan-id <value>
```

Variable definitions

Use the data in the following table to help you use the **config mlt <mlt-id> ist create ip <value> vlan-id <value>** command.

Variable	Value
ip <value>	Specifies the IP address of the peer switch
<mlt-id>	Specifies the MLT ID.
vlan-id <value>	Specifies the VLAN ID.

Procedure job aid: Example of creating an interswitch trunk from an existing MLT

The following lists the steps you need to do to create an interswitch trunk from an existing MLT.

1. Create an interswitch trunk to enable the switch.
2. Create a VLAN on the redundant aggregation switch and also on the second aggregation switch. The interswitch trunk treats the two switches as a single switch.
3. To allow the two switches to communicate, you must assign an IP address to both VLANs. The peer IP address is the IP address of the InterSwitch Trunking (IST) VLAN on the peer aggregation switch.

switch A		switch B
VLAN 20		VLAN 20
10.1.1.1. /24	<-----IST----->	10.1.1.2 /24 (see Note 1)

Note 1: Same subnet, same VLAN.

Enabling and disabling an interswitch trunk

Enable and disable an interswitch trunk.

Procedure steps

1. Enable and disable an interswitch trunk by using the following command:

```
config mlt <mlt-id> ist <enable|disable>
```

Variable definitions

Use the data in the following table to help you use the **config mlt <mlt-id> ist <enable|disable>** command.

Variable	Value
<mlt-id>	Specifies the MLT ID.

Deleting an interswitch trunk

Delete an interswitch trunk.

Prerequisites

- Before you can delete the interswitch trunk, you must disable the interswitch trunk. For more information, see [Enabling and disabling an interswitch trunk](#) on page 103.

Procedure steps

1. Delete an interswitch trunk by using the following command:

```
config mlt <mlt-id> ist delete
```

Variable definitions

Use the data in the following table to help you use the **config mlt <mlt-id> ist <enable|disable>** command.

Variable	Value
<mlt-id>	Specifies the MLT ID.

Enabling SMLT Interoperability with MSTP using CLI

Use the following procedure to enable SMLT Interoperability for MSTP.

Prerequisites

- Configure your switch to MSTP mode.
- Configure an IST MLT.

Procedure steps

1. At the system prompt, enter the following command
`config mlt <ID> ist spanning-tree enable`

Disabling SMLT Interoperability with MSTP using CLI

Use the following procedure to disable SMLT Interoperability for MSTP.

Procedure steps

1. At the system prompt, enter the following command
`config mlt <ID> ist spanning-tree disable`

**Note:**

If you want to disable SMLT Interoperability for MSTP, ensure that you remove redundant paths, for VLANs, to the existing IST path. To prevent the formation of a bridging loop and network failure, do not break or remove an MSTP device chain until you disable SMLT interoperability.

Displaying MLT IST information using CLI

Use the following procedure to display IST MLT information for your switch.

Procedure steps

1. At the prompt, enter the following command

```
show mlt ist info
```

Configuring a single port SMLT

Configure a split MultiLink trunk that uses a single port. Use single port SMLT to permit the scaling of the number of SMLT on a switch to the maximum number of available ports.

**Important:**

You cannot use SMLT on router ports. Only VLACP is supported on single port SMLT.

Prerequisites

- If the port is configured as MLT or MLT-based SMLT, you cannot configure a single port SMLT.

Procedure steps

1. Configure a single port SMLT by using the following command:

```
config ethernet <port> smlt <smltid>
```

Variable definitions

Use the data in the following table to help you use the **config ethernet <port> smlt <smltid>** command.

Variable	Value
create	Creates a single port SMLT.
delete	Deletes a single port SMLT.
info	Shows the port SMLT information.
<port>	Specifies the port number
<smltid>	Specifies the MLT ID.

Viewing the MLT port calculated by the MLT hash algorithm

View the MLT port calculated by the MLT hash algorithm to obtain information about the MLT port calculated by the distribution algorithm for R, RS, and 8800 modules.

Procedure steps

1. View the MLT port calculated by the MLT hash algorithm by using the following command:

```
config sys set hash-calc getmltindex traffic-type {non-ip|ipv4|
ipv6} dest-val <value> src-val <value> mltID <value> [src-port
<value>] [dst-port <value>]
```

Variable definitions

Use the data in the following table to help you use the **config sys set hash-calc getmltindex traffic-type <value> dest-val <value> src-val <value> mltID <value>** command.

Variable	Value
dest-val <value>	Specifies the destination IP or MAC address (string length from 1–1536). The source and destination addresses cannot have the same value.
dst-port <value>	Specifies the TCP or UDP destination port in the range 0–65535.
src-port <value>	Specifies the TCP or UDP source port in the range 0–65535.
src-val <value>	Specifies the source IP or MAC address (string length from 1–1536). The source and destination addresses cannot have the same value.
mltID <value>	Specifies the MLT ID in the range 1–256.
traffic-type {non-ip ipv4 ipv6}	Specifies the traffic type: one of non-ip, ipv4, or ipv6.

MLT and SMLT show commands

Use the MLT and SMLT show commands to display information and statistics about MLT operation in the switch.



Note:

If you disable and enable the SMLT or IST ports, the designated port information may not appear correctly.

MLT and SMLT show commands navigation

- [Viewing MLT information](#) on page 108
- [Viewing information about collision errors](#) on page 109
- [Viewing information about Ethernet errors](#) on page 109
- [Viewing MLT status](#) on page 110
- [Viewing interswitch trunk status](#) on page 110
- [Viewing SMLT status](#) on page 111
- [Viewing all ports configured for single port SMLT](#) on page 111
- [Viewing a port configured for Single Port SMLT](#) on page 111

Viewing MLT information

View all MLT information to verify all MLT information.

Procedure steps

1. View all MLT information by using the following command:

```
show mlt show-all [file <value>]
```

Variable definitions

Use the data in the following table to help you use the **show mlt show-all [file <value>]** command.

Variable	Value
<value>	Specifies the filename to which the output is directed.

Viewing information about collision errors

View information about collision errors to display information about collision errors in the specified MLT, or for all MLTs.

Procedure steps

1. View information about collision errors by using the following command:

```
show mlt error collision [<mid>]
```

Variable definitions

Use the data in the following table to help you use the **show mlt error collision [<mid>]** command.

Variable	Value
<mid>	Specifies the MLT ID

Viewing information about Ethernet errors

View information about Ethernet errors to obtain information about the types of Ethernet errors sent and received by a specified MLT or by all MLTs.



Important:

The IMAC columns refer to internal MAC address errors.

Procedure steps

1. View information about Ethernet errors by using the following command:

```
show mlt error main [ <mid> ]
```

Variable definitions

Use the data in the following table to help you use the **show mlt error main [<mid>]** command.

Variable	Value
<mid>	Specifies the MLT ID

Viewing MLT status

View the MLT status to display the status of MultiLink Trunking for the switch or for the specified MLT ID.

Procedure steps

1. View MLT status by using the following command:

```
show mlt info [<mid>]
```

Variable definitions

Use the data in the following table to help you use the **show mlt info [<mid>]** command.

Variable	Value
<mid>	Specifies the MLT ID

Viewing interswitch trunk status

View the interswitch trunk status to display the status of interswitch trunking for the switch or for the specified MLT ID.

Procedure steps

1. View interswitch trunk status by using the following command:

```
show mlt info [<mid>]
```

Variable definitions

Use the data in the following table to help you use the **show mlt info [<mid>]** command.

Variable	Value
<mid>	Specifies the MLT ID.

Viewing SMLT status

View the SMLT status to display SMLT status for the switch or for a specific SMLT ID.



Important:

The switch shows both MLT-based SMLT information and single port SMLT information.

Procedure steps

1. View SMLT status by using the following command:

```
show smlt
```

Viewing all ports configured for single port SMLT

View all ports for a single port SMLT to ensure the correct ports are configured.

Procedure steps

1. View all ports configured for single port SMLT by using the following command:

```
show ports info smlt [<vlan <value>] [port <value>]
```

Variable definitions

Use the data in the following table to help you use the **show ports info smlt** commands.

Variable	Value
<i>vlan</i> <va lue>	Specifies the VLAN ID.
<i>port</i> <va lue>	Specifies the port.

Viewing a port configured for Single Port SMLT

View a port configured for Single Port SMLT.

Procedure steps

1. View a port configured for Single Port SMLT by using the following command:

```
show ports info config [<vlan <value>] [port <value>]
```

Variable definitions

Use the data in the following table to help you use the **show ports info config** command.

Variable	Value
port <value>	Specifies a port or list of ports.
vlan <value>	Specifies a VLAN ID or list of VLAN IDs.

Chapter 9: VLACP configuration using the CLI

Configure Virtual LACP (VLACP) to implement link status control protocol at the port level. VLACP detects for end-to-end failures in the Avaya Ethernet Routing Switch 8800/8600. Note that VLACP cannot interoperate with LACP.

Job aid: Roadmap of VLACP commands

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
config ethernet <portlist> vlacp	disable
	enable
	ethertype <ethertype>
	fast-periodic-time <milliseconds>
	info
	macaddress <mac>
	slow-periodic-time <milliseconds>
	timeout {long short}
	timeout-scale <integer>
show ports info vlacp [vlan <value>] [port <value>]	
config vlacp	disable
	enable
	info
	hold-time <seconds>

Configuring VLACP on a port

Configure VLACP on a port to ensure there is end-to-end reachability. VLACP uses the Hello mechanism of LACP to periodically send Hello packets to ensure there is an end-to-end approach. When Hello packets are not received, VLACP transitions to a failure state and the port is disabled.



Important:

Changes made at the global level overrides and resets all port level settings.

Procedure steps

1. Configure VLACP on a port by using the following command:

```
config ethernet <portlist> vlacp
```

Variable definitions

Use the data in the following table to help you use the **config ethernet <portlist> vlacp** command.

Variable	Value
disable	Disables VLACP for a specific port type.
enable	Enables VLACP for a specific port type.
ethertype <integer>	Sets the VLACP protocol identification for this port. • <i>integer</i> is the ethertype value, an integer value in the range of 1 and 65535. The default value is 0x8103.
fast-periodic-time <milliseconds>	Sets the fast periodic time (in milliseconds) for a specific port type. The default value is 200.
info	Shows current level parameter settings and next level directories.
macaddress <mac>	Sets the multicast MAC address used for the VLACPDU. • <i>mac</i> is the MAC address in the following format: 0x00:0x00:0x00:0x00:0x00:0x00

Variable	Value
slow-periodic-time <milliseconds>	Sets the slow periodic time (in milliseconds) for a specific port type. The default value is 30000.
timeout {long short}	Sets the port to use the long or short timeout: • long sets the port to use the timeout-scale value multiplied by the slow-periodic-time. • short sets the port to use the timeout-scale value multiplied by the fast-periodic-time. For example, if you specify a short timeout, set the timeout-scale value to 3, and the fast-periodic-time to 400 ms, the timer expires within 1000 to 1200 ms. The default value is long.
timeout-scale <integer>	Sets a timeout scale for a specific port-type (where timeout-scale = periodic-time x timeout-scale). The default value is 3. • <i>integer</i> is the timeout scale, an integer value in the range 2 to 10.

Viewing VLACP port configuration

View the VLACP port configuration to see the port information.

Procedure steps

1. View the VLACP port configuration by using the following command:

```
show ports info vlacp [vlan <value>] [port <value>]
```

Enabling or disabling VLACP

Enable or disable the VLACP globally to reset all port level settings on the chassis.



Important:

Changes made at the global level overrides and resets all port level settings.

Procedure steps

1. To enable VLACP globally, use the following command:
`config vlacp enable`
2. To disable VLACP globally, use the following command:
`config vlacp disable`

Configuring VLACP flap detect and damping

Configure VLACP flap detect and damping to automatically shut down a VLACP link that is flapping.

Prerequisites

- Make sure that the VLACP flap interval is long enough to recognize that a port is flapping. (The default settings may not allow enough time.) A VLACP down event does not occur until the receiving side has not seen a VLACP PDU for a number of seconds (**vlacp timeout** x **vlacp timeout-scale**). Then this has to happen several times (**flap-frequency**) within a specified period of time (**flap-interval**).

To ensure that the interval is long enough, use the following formula:

flap-interval > **flap-frequency** x **vlacp timeout** x **vlacp timeout-scale**

The **flap-interval** and **flap-frequency** commands are described in this procedure.



Note:

VLACP flap detect is also triggered when there is a physical link down event on the underlying Ethernet port. This replaces the link-flap-detect global functionality, which is by default enabled.

For information on the **vlacp timeout** and **vlacp timeout-scale** commands, see [Configuring VLACP on a port](#) on page 114.

Procedure steps

To configure VLACP flap detection, use the following command:

`config ethernet <port> vlacp`

Variable definitions

The following table defines optional parameters that you enter to configure VLACP flap detect and damping.

Variable	Value
flap-detect <i><true false></i>	Enable or disable the flap detect feature. Flap detect is disabled by default.
flap-frequency <i><integer></i>	Specifies the flap frequency. The value ranges from 3-30. The default is 3.
flap-interval <i><seconds></i>	Specifies the flap interval, in seconds. The value ranges from 10-600. The default is 60.
<i><port></i>	Specifies the port or portlist.

Viewing VLACP flap detect and damping configuration

View the VLACP configuration to view the current VLACP flap detect and damping settings.

Procedure steps

1. To view VLACP configuration and VLACP Flap information, use the following command:

```
show port info vlacp
```
2. To clear VLACP Flap information, use the following command:

```
clear vlacp flap-stats [<ports>]
```

Variable definitions

The following table defines the parameters that you enter with the `show port info vlacp` command.

Variable	Value
vlacp [vlan <value>] [port <value>]	Shows VLACP port information. • <code>vlan <value></code> is the VLAN ID. • <code>port <value></code> is the port or portlist

Chapter 10: MLT and MLT with LACP CLI configuration examples

This section contains configuration examples for configuring MLT and MLT with LACP using the CLI.

For additional configuration examples about MLT with LACP and VLACP, see *Link Aggregation Control Protocol (LACP) 802.3ad and VLACP for ES and ERS Technical Configuration Guide*.

For information about SMLT configuration examples, see *Switch Clustering Split Multi-Link Trunking Technical Configuration Guide, NN48500-518*.

You can find the technical configuration guides at www.avaya.com/support with the rest of the ERS 8800/8600 documentation.

MultiLink Trunking configuration example

This configuration example shows how to create a MultiLink Trunk (MLT) and a VLAN between two Avaya Ethernet Routing Switch 8800/8600 devices. The network depicted in the following diagram shows a MLT within a VLAN used to carry user traffic.




Figure 13: MLT within a VLAN

The following procedures show how to configure switch S1 and S2 for this example.

Configuring S1:

1. Create VLAN 100:

```
ERS-8606:5#config vlan 100 create byport 1
```

2. Create multilink trunk 10:

```
ERS-8606:5#config mlt 10 create
ERS-8606:5#config mlt 10 add ports 1/1,1/2,2/1,2/2
ERS-8606:5# config mlt 10 add vlan 100
```

Configuring S2:

1. Create VLAN 100:

```
ERS-8606:5#config vlan 100 create byport 1
```

2. Create multilink trunk 10:


```
ERS-8606:5#config mlt 10 create
ERS-8606:5#config mlt 10 add ports 1/1,1/2,2/1,2/2
ERS-8606:5#config mlt 10 add vlan 100
```

Multilink with LACP configuration example


This configuration example shows how to configure and enable a MultiLink Trunk using LACP.

You must configure all aggregatable ports to use the same aggregator key used for the MLT.

The following procedures show how to configure switch S1 and S2 in the following diagram as a MLT with LACP.



Legend



11340FA

Configuring S1:

1. Create VLAN 100 and add ports to the VLAN:

```
ERS-8606:5#config vlan 100 create byport 1  
ERS-8606:5# config vlan 100 ports add 1/1-1/2,2/1-2/2
```

2. Configure LACP on S1 switch ports:

```
ERS-8606:5#config ether 1/1-1/2,2/1-2/2 lacp key 10  
ERS-8606:5# config ether 1/1-1/2,2/1-2/2 lacp aggregation true  
ERS-8606:5#config ether 1/1-1/2,2/1-2/2 lacp enable
```

3. Create MLT 10 and configure LACP. Ensure the LACP key is the same as that configured in step 2:

```
ERS-8606:5#config mlt 10 create  
ERS-8606:5#config mlt 10 lacp key 10  
ERS-8606:5#config mlt 10 lacp enable
```

Configuring S2:

1. Create VLAN 100 and add ports to the VLAN:

```
ERS-8606:5#config vlan 100 create byport 1  
ERS-8606:5#onfig vlan 100 ports add 1/1-1/2,2/1-2/2
```

2. Configure LACP on S2 switch ports:

```
ERS-8606:5#config ether 1/1-1/2,2/1-2/2 lacp key 10  
ERS-8606:5#config ether 1/1-1/2,2/1-2/2 lacp aggregation true  
ERS-8606:5#config ether 1/1-1/2,2/1-2/2 lacp enable
```

3. Create MLT 10 and configure LACP. Ensure the LACP key is the same as that configured in step 2:

```
ERS-8606:5#config mlt 10 create  
ERS-8606:5#config mlt 10 lacp key 10  
ERS-8606:5#config mlt 10 lacp enable
```


Chapter 11: Link aggregation configuration using the ACLI

MultiLink Trunking (MLT) with Link Aggregation Control Protocol (LACP) manages switch ports and port memberships to form a link aggregation group (LAG). Configure LACP to allow dynamic bundling of physical ports to form a single logical channel.

LACP can be described in terms of link aggregation operations within a single system. You can configure a single piece of equipment so it contains more than one system (from the point of view of the link aggregation operation).

Prerequisites

- Changes to LACP made at the global level overrides and resets all port level settings.



Important:

When you globally set the LACP system priority, it applies to all LACP-enabled aggregators and ports. When you enable LACP on an aggregator or port, it uses the global system priority value.

- If you make a timer change, restart LACP (globally or on the port) so the changes are consistent across the link.



Important:

Configuration changes to LACP timers are not reflected immediately. LACP timers are not reset until the next time LACP is restarted globally or on a port. This ensures consistency with peer switches.

- Standby mode for LACP aggregation groups of larger than eight ports is not supported in this release.

Job aid: Roadmap of LACP commands

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
Privileged EXEC mode	
show lacp	<cr>
	actor-admin interface [{fastethernet gigabitethernet}] [vid {vlan-id[-vlan-id][,...]}] [<portList>]
	actor-oper interface [{fastethernet gigabitethernet}] [vid {vlan-id[- vlan-id][,...]}] [<portList>]
	extension interface [{fastethernet gigabitethernet}] [vid {vlan-id[- vlan-id][,...]}] [<portList>]
	partner-admin interface [{fastethernet gigabitethernet}] [vid {vlan-id[-vlan-id][,...]}] [<portList>]
	partner-oper interface [{fastethernet gigabitethernet}] [vid {vlan-id[-vlan-id][,...]}] [<portList>]
show lacp interface	<cr>
	fastethernet [vid {vlan-id[-vlan-id] [,...]}] [<portList>]
	gigabitethernet [vid {vlan-id[-vlan- id][,...]}] [<portList>]
	mlt [<64-6399>]
	pos [vid {vlan-id[-vlan-id][,...]}] [<portList>]
Global Configuration mode	
lacp	
	aggr-wait-time <200-2000>
	enable
	fast-periodic-time <200-20000>
	slow-periodic-time <10000-30000>
	smlt-sys-id <0x00:0x00:0x00:0x00:0x00:0x00>

Command	Parameter
	system-priority <0-65535>
	timeout-scale <2-10>
Interface Configuration mode	
lacp	
	aggregation enable
	aggr-wait-time <200-2000>
	enable
	fast-periodic-time <200-20000>
	key <0-65535>
	min-links <1-8>
	mode {active, passive}
	partner-key <0-65535>
	partner-port <0-65535>
	partner-port-priority <0-65535>
	partner-state <0x00-0xff>
	partner-system-id <0x00:0x00:0x00:0x00:0x00:0x00>
	partner-system-priority <0-65535>
	priority <0-65535>
	slow-periodic-time <10000-30000>
	system-priority <0-65535>
	timeout-scale <2-10>
	timeout-time {long,short} [timeout-scale <2-10>]

Configuring global LACP parameters

Configure LACP parameters globally. When the LACP system priority is set globally, it applies to all LACP-enabled aggregators and ports. When LACP is enabled on an aggregator or a port, it will use the global system priority value.

New global parameter configuration is in effect when LACP is restarted globally or per port.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

**Important:**

Changes made at the global level override and reset all port level settings.

Procedure steps

1. Configure LACP globally by using the following command:

```
lacp
```

Variable definitions

The following table defines parameters that you enter after the **lacp** command.

Variable	Value
aggr-wait-time <200–2000>	Sets the aggregation wait time (in milliseconds) globally. The default value is 2000.
enable	Enables LACP globally. To disable LACP globally, use the no operator to remove this configuration. To set this option to the default value, use the default operator with the command.
fast-periodic-time <200–20000>	Sets the fast-periodic time (in milliseconds) globally. The default value is 1000.
slow-periodic-time <10000–30000>	Sets the slow periodic time globally. The default value is 30000.
smlt-sys-id <0x00:0x00:0x00:0x00:0x00:0x00>	Sets the LACP system ID globally. Enter a MAC address in the following format: 0x00:0x00:0x00:0x00:0x00:0x00.
system-priority <0-65535>	Sets the global LACP system priority. The default value is 32768.
timeout-scale <2-10>	Sets the timeout scale globally. The default value is 3. To set this option to the default value, use the default operator with the command.

Configuring LACP on a port

Configure LACP on a port to enable or disable LACP on the selected ports.

A port can operate in active or passive mode. You can configure a port to be an individual link or an aggregated link.

**Important:**

Changes made at the global level override and reset all port level settings.

Prerequisites

- You must log on to the Interface Configuration mode in the ACLI.

Procedure steps

1. Configure LACP on a port by using the following command:

```
lacp enable
```

Variable definitions

The following table defines optional parameters that you enter after the **lacp enable** command.

Variable	Value
aggregation enable	Sets the port as aggregatable. Use the no operator to remove this configuration. To set this option to the default value, use the default operator with the command. The default value is disable.
aggr-wait-time <200–2000>	Sets the aggregation wait time (in milliseconds) for this port. The default value is 2000 ms.
fast-periodic-time <200–20000>	Sets the fast-periodic time (in milliseconds) for this port. The default is 1000 ms.
key <0–65535>	Sets the aggregation key for this port.

Variable	Value
mode {active, passive}	Sets the LACP mode to be active or passive. The default value is active.
partner-key <0–65535>	Sets the partner administrative key. The default value is 0.
partner-port <0–65535>	Sets the partner administrative port value. The default value is 0.
partner-port-priority <0–65535>	Sets the partner administrative port priority value. The default value is 0.
partner-state <0x00–0xff>	Sets the partner administrative state bitmask. Specify the partner administrative state bitmap in the range 0x0–0xff. The bit to state mapping is Exp, Def, Dis, Col, Syn, Agg, Time, and Act. For example, to set the two partner-state parameters • Act = true • Agg = true specify a value of 0x05 (bitmap = 00000101).
partner-system-id <0x00:0x00:0x00:0x00:0x00:0x00>	Sets the partner administrative system ID. Specify a MAC address in the format 0x00:0x00:0x00:0x00:0x00:0x00.
partner-system-priority <0–65535>	Sets the partner administrative system priority value. The default value is 0.
priority <0–65535>	Sets the port priority. The default value is 32768. To set this option to the default value, use the default operator with the command.
slow-periodic-time <10000-30000>	Sets the slow periodic time for this port. The default is 30000 ms. To set this option to the default value, use the default operator with the command.
system-priority <0-65535>	Sets the system priority for this port. The default value is 32768.
timeout-scale <2-10>	Sets a timeout scale for this port. The default value is 3. To set this option to the default value, use the default operator with the command.
timeout-time {long,short}	Sets the timeout to either long or short. To set this option to the default value, use the default operator with the command. The default value is long.

Viewing LACP configuration information

View LACP configuration information to determine the LACP parameters and to ensure your configuration is correct.

Prerequisites

- You must log on to the Privileged EXEC mode in the ACLI.

Procedure steps

1. View LACP configuration information by using the following command:

```
show lacp
```

Variable definitions

The following table defines optional parameters that you enter after the **show lacp** command.

Variable	Value
<cr>	Shows LACP global configuration information.
actor-admin interface [{fastethernet gigabitethernet}] [vid {vlan-id[-vlan-id][,...]}] [<portList>]	Shows LACP actor administrative information for all interfaces or the specified interface. • vlan-id is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4094. • portList is the port or port list.
actor-oper interface [{fastethernet gigabitethernet}] [vid {vlan-id[-vlan-id][,...]}] [<portList>]	Shows LACP actor operational information for all interfaces or the specified interface. • vlan-id is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4094. • portList is the port or port list.

Variable	Value
extension interface [{fastethernet gigabitethernet}] [vid {vlan-id[-vlan-id][,...]}] [<portList>]	Shows LACP timer information for all interfaces or the specified interface. • vlan-id is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4094. • portList is the port or port list.
interface [{fastethernet gigabitethernet pos}] [vid {vlan-id[-vlan-id][,...]}] [<portList>]	Shows all LACP port configuration information for all interfaces or the specified interface.
interface mlt [<64–6399>]	Shows the MLT LACP information for all MLTs or the specific MLT index.
partner-admin interface [{fastethernet gigabitethernet}] [vid {vlan-id[-vlan-id][,...]}] [<portList>]	Shows LACP partner administrative information for all interfaces or the specified interface. • vlan-id is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4094. • portList is the port or port list.
partner-oper interface [{fastethernet gigabitethernet}] [vid {vlan-id[-vlan-id][,...]}] [<portList>]	Shows LACP partner operational information for all interfaces or the specified interface. • vlan-id is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4094. • portList is the port or port list.

Chapter 12: MLT and SMLT link aggregation configuration using the ACLI

Configure link aggregation to provide link level redundancy and increase load sharing. Multi-Link Trunking (MLT) is a link aggregation technology that allows grouping several physical Ethernet links into one logical Ethernet link to provide fault-tolerance and high-speed links between routers, switches, and servers. SMLT is an option that improves Layer 2 (bridged) resiliency.

Job aid: Roadmap of MLT and SMLT commands

The following command lists the command and their parameters that you use to complete the procedures in this section.

Command	Parameter
<i>Privileged EXEC mode</i>	
show mlt	
show mlt error	
	collision [<1-256>]
	main [<1-256>]
show mlt stats <mlt ID>	
<i>Global Configuration mode</i>	
mlt <mlt ID>	enable name <0-20>
	encapsulation dot1q
	member <portList>
	name <0-20>
	ntstg enable
	vlan <vid>
show smlt	<cr>
	{fastethernet gigabitethernet} [<portList>]

Command	Parameter
show smlt mlt	
<i>MLT Interface Configuration mode</i>	
ist peer-ip default ist peer-ip	
ist peer-ip <A.B.C.D> vlan <vlan-id>	
smlt <port ID>	
smlt <smlt ID>	

Configuring an MLT

Configure an MLT to set up MLTs on the switch.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Configure an MLT by using the following command:

```
mlt <mlt ID>
```

Variable definitions

Use the data in the following table to help you use the **mlt <mlt ID>** command.

Variable	Value
enable	Creates and enables a new MLT.
encapsulation dot1q	Sets encapsulation. dot1q enables trunking on the MLT.
member <portList>	Adds ports to this MLT.
<mlt ID>	Specifies the MLT ID in the range of 1–256.
name <0-20>	Changes the name for this MLT in the range of 0–20 characters.

Variable	Value
ntstg enable	Enables NTSTG. The default value is true.
vlan <1-4094>	Specifies a VLAN ID to add to this MLT.

Adding ports to an MLT LAG

Add ports to an MLT LAG to add an existing VLAN to a link aggregation configuration.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Adding ports to an MLT LAG by using the following command:

```
mlt <1-32> member <portlist>
```

Removing ports from an MLT LAG

Remove ports from an MLT LAG.

Prerequisites

- Before removing a port member from an MLT, you need to first disable the port. This ensures that the other side will bring its corresponding port member down. This achieves parity on both sides and avoids traffic disruptions.
- You must log on to the Global configuration mode in the ACLI.

Procedure steps

1. Remove ports from an MLT LAG by using the following command.

```
no mlt <mlt-id> member <portlist>
```

Configuring LACP on an MLT

Configure an MLT with LACP to use the dynamic link aggregation function.



Important:

Attach ports to an aggregator only if their system priorities are the same; otherwise, they are considered to be operating in two different switches. You can attach ports to an aggregator only if their keys are the same.

Prerequisites

- You must log on to the Global configuration mode in the ACLI.

Procedure steps

1. To log on to the Interface MLT Configuration mode for the MLT that you want configure LACP on enter the following at the ACLI Global configuration command prompt:

```
interface mlt <mltid>
```



Tip:

Specify an MLT ID in the range 1-256 for <mltid>.

2. Configure LACP on an MLT by using the following command:

```
lacp enable key <integer>
```

Variable definitions

Use the data in the following table to help you use the **lacp enable key <integer>** command.

Variables	Value
key <integer>	Sets LACP aggregator key for a specific MLT. • <i>integer</i> is the LACP actor admin key. The default value is 0.

Variables	Value
min-links <1-8>	Configures the minimum number of active links (MinLink) required for the LAG to remain in forwarding state. If the number of active links falls below this number, the Avaya Ethernet Routing Switch 8800/8600 declares the entire LAG down. The default value is 1.
system-priority <integer>	Sets the LACP system priority for a specific MLT. • <i>integer</i> is the system priority in the range 0–65535. The default value is 32768.

Creating a SMLT from an existing MLT

Create a SMLT from an existing MLT to split physical ports between two switches to improve resiliency and provide active load sharing.

Prerequisites

- Create an MLT before you create a split in the MLT.
- You must log on to the Global configuration mode in the ACLI.

Procedure steps

1. To log on to the Interface MLT Configuration mode for the MLT that you want to split enter the following at the ACLI Global configuration command prompt:

```
interface mlt <mltid>
```



Tip:

Specify an MLT ID in the range 1-256 for <mltid>.

2. Create a SMLT from an existing MLT by using the following command:

```
smlt <smltid>
```

Variable definitions

Use the data in the following table to help you use the **smlt <smltid>** command.

Variable	Value
<smltid>	Specifies the Split MLT ID in the range of 1–256. The value must match the peer switch SMLT ID.

Creating an interswitch trunk

Create an interswitch trunk from an existing MLT.

Prerequisites

- You must log on to the Global configuration mode in the ACLI.

Procedure steps

1. To log on to the Interface MLT Configuration mode for the MLT that you want to configure enter the following at the Global configuration command prompt:

```
interface mlt <mltid>
```



Tip:

Specify an MLT ID in the range 1-256 for <mltid>.

2. Create an interswitch trunk by using the following command:

```
ist peer-ip <A.B.C.D> vlan <vid>
```

Variable definitions

Use the data in the following table to help you use the **ist peer-ip** command.

Variable	Value
<A.B.C.D>	Specify the peer IP address—the IP address of the IST VLAN on the other aggregation switch.
<1–4094>	Specify the VLAN ID.

Enabling an interswitch trunk

Enable an interswitch trunk.

Prerequisites

- You must log on to the Global configuration mode in the ACLI.

Procedure steps

1. To log on to the Interface MLT Configuration mode for the MLT that you want to configure enter the following at the Global configuration command prompt:

```
interface mlt <mltid>
```



Tip:

Specify an MLT ID in the range 1–256 for <mltid>.

2. Enable an interswitch trunk by using the following command:

```
ist enable
```

3. To disable an interswitch trunk, enter the following command:

```
no ist enable
```

Deleting an interswitch trunk

Delete an interswitch trunk.

Prerequisites

- Before you can delete the interswitch trunk, you must disable the interswitch trunk. For more information, see [Enabling and disabling an interswitch trunk](#) on page 103.
- You must log on to the MLT Interface mode in the ACLI.

Procedure steps

1. Delete an interswitch trunk by using the following command:

```
no ist peer-ip
```

Job aid: Example of deleting an interswitch trunk

The following example shows the commands that you use to delete an interswitch trunk.

Procedure steps

1. From Global configuration mode, log into the MLT interface:

```
ERS-8600:6(config)# interface mlt 2
```

2. Disable the IST:

```
ERS-8600:6(config)# no ist enable
```

3. Delete the IST:

```
ERS-8600:6(config)# no ist peer-ip
```

Configuring a single port SMLT

Configure a split MultiLink trunk that uses a single port. Use single port SMLT to permit the scaling of the number of SMLT on a switch to the maximum number of available ports.



Important:

You cannot use SMLT on router ports. Only VLACP is supported on single port SMLT.

Prerequisites

- If the port is configured as MLT or MLT-based SMLT, you cannot configure a single port SMLT.
- You must log on to the Gigabit Interface Configuration mode in the ACLI.

Procedure steps

1. Configure a single port SMLT by using the following command:

```
smlt <smltid>
```

Variable definitions

Use the data in the following table to help you use the **smlt <smltid>** command.

Variable	Value
<smltid>	Specifies the Split MLT ID. The value ranges from 1–512.

Enabling SMLT Interoperability with MSTP using ACLI

Use the following procedure to enable SMLT Interoperability for MSTP.

Prerequisites

- Configure your switch to MSTP mode.
- Configure an IST MLT.

Procedure steps

1. At the prompt, enter the following commands:

```
ist mlt <ID>
```

```
ist spanning-tree
```

Disabling SMLT Interoperability with MSTP using ACLI

Use the following procedure to disable SMLT Interoperability for MSTP.

Procedure steps

1. At the prompt, enter the following commands:

```
ist mlt <ID>  
no ist spanning-tree
```



Note:

If you want to disable SMLT Interoperability for MSTP, ensure that you remove redundant paths, for VLANs, to the existing IST path. To prevent the formation of a bridging loop and network failure, do not break or remove an MSTP device chain until you disable SMLT interoperability.

Displaying MLT IST information using ACLI

Use the following procedure to display MLT IST information.

Prerequisites

- Log on to the Privileged Exec mode

Procedure steps

1. At the prompt, enter the following command:

```
show mlt ist
```

Viewing the MLT port calculated by the MLT hash algorithm

View the MLT port calculated by the MLT hash algorithm to obtain information about the MLT port calculated by the distribution algorithm for R, RS, and 8800 modules.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

View the MLT port calculated by the MLT hash algorithm to view the MLT port calculated by the distribution algorithm for R, RS, and 8800 modules.

Procedure steps

1. View the MLT port calculated by the MLT hash algorithm by using the following command:

```
sys hash-calc getmltindex traffic-type <ipv4|ipv6|non-ip> dest-
val <WORD 1-1536> src-val <WORD 1-1536> mltID <1-256> [dest-
port <0-65535>] [src-port <0-65535>]
```

The source and destination addresses cannot have the same value.

Variable definitions

Use the data in the following table to help you use the **sys hash-calc** command.

Variable	Value
dest-val <value>	Specifies the destination address in the range from 1–1536. The source and destination addresses cannot have the same value.
src-val <value>	Specifies the source address in the range from 1–1536. The source and destination addresses cannot have the same value.
mltID <value>	Specifies the MLT ID. The value ranges from 1–256.
traffic-type {ipv4 ipv6 non-ip}	The type of traffic. Specify one of ipv4, ipv6, or non-ip.

MLT and SMLT show commands

This section describes the MLT and SMLT show commands. Use the commands to display information and statistics about MLT operation in the switch commands.



Note:

If you disable and enable the SMLT or IST ports, the designated port information may not appear correctly.

MLT and SMLT show commands navigation

- [Viewing information about collision errors](#) on page 142
- [Viewing information about Ethernet errors](#) on page 143
- [Viewing all ports configured for single port SMLT](#) on page 143

Viewing information about collision errors

View information about collision errors to obtain information about collision errors in the specified MLT, or for all MLTs.

Prerequisites

- You must log on to the Privileged EXEC mode in the ACLI.

Procedure steps

1. View information about collision errors by using the following command:

```
show mlt error collision [<mltid>]
```

Variable definitions

Use the data in the following table to help you use the **show mlt error collision** command.

Variable	Value
<mltid>	Specifies the MLT ID. The value ranges from 1–256.

Viewing information about Ethernet errors

View information about Ethernet errors to display information about the types of Ethernet errors sent and received by the specified MLT or all MLTs.

**Important:**

The IMAC columns refer to internal MAC address errors.

Prerequisites

- You must log on to the Privileged EXEC mode in the ACLI.

Procedure steps

1. View information about Ethernet errors by using the following command:

```
show mlt error main [<mltid>]
```

Variable definitions

Use the data in the following table to help you use the **show mlt error main** command.

Variable	Value
<mltid>	Specifies the MLT ID. The value ranges from 1–256.

Viewing all ports configured for single port SMLT

View all ports for a single port SMLT to ensure the correct ports are configured.

Prerequisites

- You must log on to the Privileged EXEC mode in the ACLI.

Procedure steps

1. View all ports configured for single port SMLT by using the following command:

```
show smlt {fastethernet|gigabitethernet|mlt|pos}
```

Variable definitions

Use the data in the following table to help you use the **{fastethernet|gigabitethernet|mlt|pos}** command.

Variable	Value
{fastethernet gigabitethernet pos} [<portList>]	Displays SMLT information for the interface and ports that you specify. • portList is the port or port list.
mlt	Displays SMLT information for the MLT interface.

Chapter 13: VLACP configuration using the ACLI

Configure Virtual LACP (VLACP) to implement link status control protocol at the port level. VLACP detects for end-to-end failures in the Avaya Ethernet Routing Switch 8800/8600. Note that VLACP cannot interoperate with LACP.

Job aid: Roadmap of VLACP commands

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
Privileged EXEC mode	
show vlacp interface	<cr>
	fastethernet [vid {vlan-id[-vlan-id][,...]}] [<portList>]
	gigabitethernet [vid {vlan-id[-vlan-id][,...]}] [<portList>]
Global Configuration mode	
vlacp enable	
no vlacp enable	
Interface Configuration mode	
vlacp	enable
	ethertype <1-65535>
	fast-periodic-time <10-20000>
	funcmac-addr <0x00:0x00:0x00:0x00:0x00:0x00>

Command	Parameter
	slow-periodic-time <10000-30000>
	timeout {long short}
	timeout-scale <2-10>

Configuring VLACP on a port

Configure VLACP on a port to ensure there is end-to-end reachability. VLACP uses the Hello mechanism of LACP to periodically send Hello packets to ensure there is an end-to-end approach. When Hello packets are not received, VLACP transitions to a failure state and the port is disabled.



Important:

Changes made at the global level override and reset all port level settings.

Prerequisites

You must log on to the Interface Configuration mode in the ACLI.

Procedure steps

1. Configure VLACP on a port by using the following command:

```
vlacp
```

Variable definitions

The following table defines parameters that you enter after the **vlacp** command.

Variable	Value
enable	Enables VLACP for this port. The default value is disable.
ethertype <1-65535>	Sets the VLACP protocol identification for this port.
fast-periodic-time <10-20000>	Sets the fast periodic time (in milliseconds) for this port. The default value is 200.

Variable	Value
funcmac-addr <0x00:0x00:0x00:0x00:0x00:0x00>	Sets the multicast MAC address used for the VLACPDU. Specify a MAC address in the format 0x00:0x00:0x00:0x00:0x00:0x00.
slow-periodic-time <10000-30000>	Sets the slow periodic time (in milliseconds) for a specific port type. The default value is 30000.
timeout {long short}	Sets the port to use the long or short timeout: • long sets the port to use the timeout-scale value multiplied by the slow-periodic-time. • short sets the port to use the timeout-scale value multiplied by the fast-periodic-time. For example, if you specify a short timeout, set the timeout-scale value to 3, and the fast-periodic-time to 400 ms, the timer will expire within 1000 to 1200 ms. To set this option to the default value (long), use the default operator with the command.
timeout-scale <2-10>	Sets a timeout scale for this port used to calculate the timeout. The default value is 3. To set this option to the default value, use the default operator with the command.

Viewing the VLACP port configuration

View the VLACP port configuration to show the port VLACP configuration.

Prerequisites

- You must log on to the Privileged EXEC mode in the ACLI.

Procedure steps

1. View the VLACP port configuration for all interfaces by using the following command:

```
show vlacp interface
```

Variable definitions

The following table defines optional parameters that you enter after the **show vlacp interface** command.

Variable	Value
fastethernet [vid {vlan-id[-vlan-id][,...]}] [<portList>]	Displays the VLACP configuration for the Fastethernet port interface. • vlan-id is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4094. • portList is the port or port list.
gigabitethernet [vid {vlan-id[-vlan-id][,...]}] [<portList>]	Displays the VLACP configuration for the GigabitEthernet port interface. • vlan-id is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4094. • portList is the port or port list.

Enabling or disabling VLACP globally

Enable or disable the VLACP globally to reset all port level settings on the chassis.



Important:

Changes you make at the global level override and reset all port level settings.

Prerequisites

You must log on to the Global configuration mode in the ACLI.

Procedure steps

1. To globally enable VLACP, use the following command:

```
vlacp enable
```
2. To globally disable VLACP, use the following command:

```
no vlacp enable
```

Configuring VLACP flap detect and damping

Configure VLACP flap detect and damping to automatically shut down a VLACP link that is flapping.

Prerequisites

- You must log on to the Interface Configuration mode in the ACLI.
- Make sure that the VLACP flap interval is long enough to recognize that a port is flapping. (The default settings may not allow enough time.) A VLACP down event does not occur until the receiving side has not seen a VLACP PDU for a number of seconds (**vlacp timeout** x **vlacp timeout-scale**). Then this has to happen several times (**flap-frequency**) within a specified period of time (**flap-interval**).

To ensure that the interval is long enough, use the following formula:

flap-interval > **flap-frequency** x **vlacp timeout** x **vlacp timeout-scale**

The **flap-interval** and **flap-frequency** commands are described in this procedure.



Note:

VLACP flap detect is also triggered when there is a physical link down event on the underlying Ethernet port. This replaces the link-flap-detect global functionality, which is by default enabled.

For information on the **vlacp timeout** and **vlacp timeout-scale** commands, see [Configuring VLACP on a port](#) on page 146.

Procedure steps

1. To enable VLACP flap detection, use the following command:
`vlacp flap-detect enable`
2. To define the VLACP flap frequency, use the following command:
`vlacp flap-frequency <frequency>`
3. To define the VLACP flap interval, use the following command:

```
vlacp flap-interval <interval>
```

Variable definitions

The following table defines optional parameters that you enter to configure VLACP flap detect and damping.

Variable	Value
enable	Turns on VLACP flap detect and damping. The feature is disabled by default. To disable the feature, use the <code>no</code> operator with the command
<frequency>	Specifies the number of flaps, per interval. The value ranges from 3-30. The default value is 3. To set the default value, use the <code>default</code> operator with the command
<interval>	Specifies the flap interval, in seconds. The value ranges from 10-600. The default value is 60. To set the default value, use the <code>default</code> operator with the command

Viewing VLACP flap detect and damping configuration

View the VLACP configuration to view the current flap detect and damping settings.

Procedure steps

1. To view VLACP configuration and VLACP Flap information, use the following command:

```
show vlacp
```
2. To clear VLACP information, use the following command:

```
clear vlacp flap-stats port
```

Variable definitions

The following table defines the parameters that you enter with the `show vlacp` command.

Variable	Value
<code>interface gigabitethernet <port> vid <vid></code>	Shows the current configured value for VLACP at the interface. • <i><port></i> is the port or port range. • <i><vid></i> is the VLAN ID.

Chapter 14: Customer Service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com or go to one of the pages listed in the following sections.

Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

Getting product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

Getting help from a distributor or reseller

If you purchased a service contract for your Avaya product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.

Glossary

Address Resolution Protocol (ARP)	Maps an IP address to a physical machine address, for example, maps an IP address to an Ethernet media access control (MAC) address.
Border Gateway Protocol (BGP)	An inter-domain routing protocol that provides loop-free inter-domain routing between Autonomous Systems (AS) or within an AS.
Distributed MultiLink Trunking (DMLT)	A point-to-point connection that aggregates similar ports from different modules to logically act like a single port, but with the aggregated bandwidth.
Internet Protocol multicast (IPMC)	The technology foundation for audio and video streaming, push applications, software distribution, multipoint conferencing, and proxy and caching solutions.
Layer 2	The Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
Link Aggregation Control Protocol (LACP)	A protocol that exists between two endpoints to bundle links into an aggregated link group for bandwidth increase and link redundancy.
Link Aggregation Control Protocol Data Units (LACPDU)	Link aggregation control protocol data unit (LACPDU) is used for exchanging information among LACP-enabled devices.
link aggregation group (LAG)	A group that increases the link speed beyond the limits of one single cable or port, and also increases the redundancy for higher availability.
Media Access Control (MAC)	Arbitrates access to and from a shared medium.
MultiLink Trunking (MLT)	A method of link aggregation that uses multiple Ethernet trunks aggregated to provide a single logical trunk. A multilink trunk provides the combined bandwidth of multiple links and the physical layer protection against the failure of a single link.
Multiple Spanning Tree Protocol (MSTP)	Configures multiple instances of the Rapid Spanning Tree Protocol (RSTP) on the switch.

Open Shortest Path First (OSPF)	A link-state routing protocol used as an Interior Gateway Protocol (IGP).
Protocol Independent Multicast, Sparse Mode (PIM-SM)	Adds a Rendezvous Point router to avoid multicast-data flooding. Use PIM-SM when receivers for multicast data are sparsely distributed throughout the network.
Rapid Spanning Tree Protocol (RSTP)	Reduces the recovery time after a network breakdown. RSTP enhances switch-generated Topology Change Notification (TCN) packets to reduce network flooding.
Simple Loop Prevention Protocol (SLPP)	Simple Hello Protocol that prevents loops in a Layer 2 network (VLAN).
spanning tree	A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.
Split MultiLink Trunking (SMLT)	An Avaya extension to IEEE 802.1AX (link aggregation), provides nodal and link failure protection and flexible bandwidth scaling to improve on the level of Layer 2 resiliency.
Virtual Link Aggregation Control Protocol (VLACP)	A Layer 2 handshaking protocol that can detect end-to-end failure between two physical Ethernet interfaces.
virtual router (VR)	An abstract object managed by the Virtual Router Redundancy Protocol (VRRP) that acts as a default router for hosts on a shared LAN.
Virtual Router Redundancy Protocol (VRRP)	A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place.