



Configuring Link Aggregation, MLT, and SMLT on Avaya Virtual Services Platform 9000

Release 4.1
NN46250-503
Issue 07.01
October 2015

© 2011-2015, Avaya Inc.
All Rights Reserved.

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 may generally make available to users of its products and Hosted Services. 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 websites 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 Avaya hardware and software. 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 website: <http://support.avaya.com> or such successor site as designated by Avaya. Please note that if You acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to You by said Avaya Channel Partner and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, [HTTP://SUPPORT.AVAYA.COM/LICENSEINFO](http://support.avaya.com/licenseinfo) OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AVAYA CHANNEL PARTNER. 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 CHANNEL PARTNER; 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").

Avaya grants You a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The

applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. "Software" means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. "Designated Processor" means a single stand-alone computing device. "Server" means a Designated Processor that hosts a software application to be accessed by multiple users. "Instance" means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine ("VM") or similar deployment.

Licence types

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only on a number of Designated Processors up to the number indicated in the order. Avaya may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Avaya through electronic means established by Avaya specifically for this purpose.

Heritage Nortel Software

"Heritage Nortel Software" means the software that was acquired by Avaya as part of its purchase of the Nortel Enterprise Solutions Business in December 2009. The Heritage Nortel Software currently available for license from Avaya is the software contained within the list of Heritage Nortel Products located at <http://support.avaya.com/LicenseInfo> under the link "Heritage Nortel Products" or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants You a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or for communication with Avaya equipment. Charges for Heritage Nortel Software may be based on extent of activation or use authorized as specified in an order or invoice.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Avaya. All content on this site, the documentation, Hosted Service, 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

"Third Party Components" mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements ("Third Party Components"), which contain terms regarding the rights to use certain portions of the Software ("Third Party Terms"). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Avaya's website at: <http://support.avaya.com/Copyright> or such successor site as designated by Avaya. You agree to the Third Party Terms for any such Third Party Components.

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 website: <http://support.avaya.com> or such successor site as designated by Avaya. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: <http://support.avaya.com>, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: <http://support.avaya.com> for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: <http://support.avaya.com> (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, Hosted Service(s), 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, Hosted Service(s) 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. Linux[®] is the registered trademark of Linus Torvalds in the U.S. and other countries.

Contents

Chapter 1: Introduction	6
Purpose.....	6
Related resources.....	6
Documentation.....	6
Training.....	6
Viewing Avaya Mentor videos.....	6
Support.....	7
Searching a documentation collection.....	7
Chapter 2: New in this release	9
Features.....	9
Other changes.....	9
Chapter 3: Link Aggregation Control Protocol	10
Link aggregation overview.....	10
LACP configuration considerations.....	11
LACP configuration using ACLI.....	12
Configuring global LACP parameters.....	13
Configuring LACP on a port.....	14
Viewing LACP configuration information.....	17
LACP configuration using EDM.....	20
Configuring global LACP parameters.....	20
Configuring LACP parameters.....	22
Configuring LACP on a port.....	23
Chapter 4: MultiLink Trunking and Split MultiLink Trunking	29
Link aggregation overview.....	29
MultiLink Trunking.....	30
MultiLink Trunking with LACP.....	33
Split MultiLink Trunking.....	35
MLT and SMLT configuration considerations.....	43
MLT and SMLT link aggregation configuration using the ACLI.....	46
Configuring an MLT.....	46
Adding ports to an MLT LAG.....	48
Removing ports from an MLT LAG.....	48
Configuring LACP on an MLT.....	49
Creating an SMLT from an existing MLT.....	50
Creating an interswitch trunk.....	51
Enabling an interswitch trunk.....	52
Deleting an interswitch trunk.....	53
Viewing the MLT port calculated by the MLT hash algorithm.....	53
Viewing information about collision errors.....	54

Viewing information about Ethernet errors.....	55
Viewing all ports configured for SMLT.....	55
MLT and SMLT link aggregation configuration using EDM.....	56
Adding a multilink, interswitch, or LACP trunk.....	56
Adding ports to an MLT.....	59
Viewing trunks.....	59
Configuring an IST MLT.....	60
Editing an IST.....	61
MLT configuration examples.....	62
MultiLink Trunking.....	62
MultiLink Trunking with Link Aggregation Control Protocol.....	62
MLT network topology and configuration reference.....	63
Example 1: Switch-to-switch MLT.....	64
Example 2: Switch-to-server MLT.....	64
Example 3: Client/Server MLT.....	65
Chapter 5: Virtual Link Aggregation Control Protocol.....	67
Virtual Link Aggregation Control Protocol.....	67
VLACP configuration using ACLI.....	69
Configuring VLACP on a port.....	69
Viewing the VLACP port configuration.....	71
Enabling or disabling VLACP globally.....	72
VLACP configuration using EDM.....	73
Enabling VLACP globally.....	73
Configuring VLACP on a port.....	73
Glossary.....	75

Chapter 1: Introduction

Purpose

This document contains conceptual and procedural information to help you configure and manage link aggregation and MultiLink Trunking on the Avaya Virtual Services Platform 9000. This document also provides instructions for using both Enterprise Device Manager (EDM) and Avaya Command Line Interface (ACLI).

Related resources

Documentation

See *Documentation Reference for Avaya Virtual Services Platform 9000*, NN46250-100 for a list of the documentation for this product.

Training

Ongoing product training is available. For more information or to register, you can access the website at <http://avaya-learning.com/>.

Course code	Course title
4D00010E	Knowledge Access: ACIS - Avaya ERS 8000 and VSP 9000 Implementation
5D00040E	Knowledge Access: ACSS - Avaya VSP 9000 Support

Viewing Avaya Mentor videos

Avaya Mentor videos provide technical content on how to install, configure, and troubleshoot Avaya products.

About this task

Videos are available on the Avaya Support website, listed under the video document type, and on the Avaya-run channel on YouTube.

Procedure

- To find videos on the Avaya Support website, go to <http://support.avaya.com> and perform one of the following actions:
 - In **Search**, type `Avaya Mentor Videos` to see a list of the available videos.
 - In **Search**, type the product name. On the Search Results page, select **Video** in the **Content Type** column on the left.
- To find the Avaya Mentor videos on YouTube, go to www.youtube.com/AvayaMentor and perform one of the following actions:
 - Enter a key word or key words in the **Search Channel** to search for a specific product or topic.
 - Scroll down Playlists, and click the name of a topic to see the available list of videos posted on the website.

Note:

Videos are not available for all products.

Support

Go to the Avaya Support website at <http://support.avaya.com> for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Searching a documentation collection

On the Avaya Support website, you can download the documentation library for a specific product and software release to perform searches across an entire document collection. For example, you can perform a single, simultaneous search across the collection to quickly find all occurrences of a particular feature. Use this procedure to perform an index search of your documentation collection.

Before you begin

- Download the documentation collection zip file to your local computer.
- You must have Adobe Acrobat or Adobe Reader installed on your computer.

Procedure

1. Extract the document collection zip file into a folder.

2. Navigate to the folder that contains the extracted files and open the file named `<product_name_release>.pdx`.
3. In the Search dialog box, select the option **In the index named `<product_name_release>.pdx`**.
4. Enter a search word or phrase.
5. Select any of the following to narrow your search:
 - Whole Words Only
 - Case-Sensitive
 - Include Bookmarks
 - Include Comments
6. Click **Search**.

The search results show the number of documents and instances found. You can sort the search results by Relevance Ranking, Date Modified, Filename, or Location. The default is Relevance Ranking.

Chapter 2: New in this release

The following sections detail what is new in *Configuring Link Aggregation, MLT, and SMLT on Avaya Virtual Services Platform 9000*, NN46250-503 for Release 4.1.

Features

See the following section for feature-related updates:

VLACP sequence-num

Release 4.1 supports the ability to enable or disable Virtual Link Aggregation Control Protocol (VLACP) sequence numbers. For more information, see [Configuring VLACP on a port](#) on page 69.

Other changes

There are no other changes.

Chapter 3: Link Aggregation Control Protocol

This section provides the concepts and procedures you need to configure the Link Aggregation Control Protocol (LACP) to dynamically aggregate links as they become available to a trunk group.

Link aggregation overview

Link aggregation provides link level redundancy and increases load sharing. Use Link aggregation to bundle the ports into a port group, which is represented as one logical interface to the Media Access Control (MAC) layer.

The Avaya Virtual Services Platform 9000 supports the following types of link aggregation:

- MultiLink Trunking (MLT)—a statically configured link bundling method. MLT is not standards based, but it interoperates with static link methods of other vendors. If you aggregate links from multiple modules together, MLT can provide not only port redundancy, but also module redundancy.
- IEEE 802.3ad based link aggregation, through the Link Aggregation Control Protocol (LACP), dynamically aggregates links as they become available to a trunk group. Link Aggregation Control Protocol dynamically detects whether links can be aggregated into a link aggregation group (LAG) and does so after links become available. Link Aggregation Control Protocol also provides link integrity checking at Layer 2 for all links within the LAG.

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

The Virtual Services Platform 9000 software offers LACP functionality layered with MLT. This document uses the term MLT with LACP to refer to this functionality.

An Interswitch Trunk (IST) is a special form of aggregate link. ISTs use a link aggregation mechanism such as a Multilink Trunk (MLT) to create a single logical switch out of two physical devices. Creating one logical switch allows the devices to share resources, enhance fault tolerance, and create network redundancy.

Split MultiLink Trunking (SMLT) is an option that improves Layer 2 (bridged) resiliency by providing for the addition of a Virtual Services Platform 9000 failure redundancy with subsecond failover, on top of all standard MLT link failure protection and flexible bandwidth scaling functionality. Use SMLT to connect a device that supports some form of link aggregation, be it a switch or a server, to two

distinct separate SMLT endpoints or switches. These SMLT devices form a Switch Cluster and are referred to as an Interswitch Trunk (IST) Core Switch pair.

You can also use LACP on SMLT configurations. The Virtual Services Platform 9000 provides modifications to the LACP in SMLT configurations. LACP-capable devices can connect to an SMLT aggregation pair. Avaya recommends that you do not configure LACP on the IST MLT.

Virtual LACP (VLACP) is an Avaya modification that provides end-to-end failure detection. The VLACP is not a link aggregation protocol; VLACP implements link status control protocol at the port level. This mechanism periodically checks 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 an MLT. Avaya recommends that you do not configure VLACP on LACP-enabled ports. Virtual LACP does not operate properly with LACP. You can configure VLACP with an SMLT configuration.

LACP configuration considerations

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

When you configure LACP on the IST Core Switch pair, you must configure the same LACP `smlt-sys-id` on each switch in the pair.

LACP priority

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

- Port priority—determines which ports are aggregated into a LAG that has more than eight ports configured to it, as in a standby-port configuration.
- System priority—generates the switch ID after communicating with other systems. 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

You can use the LACP keys to determine which ports are eligible for link aggregation. The LACP keys are defined by the ports after you configure the multilink trunk. You can aggregate the ports key that match the MLT key into that multilink trunk.

- Keys do not have to match between two LACP peers.
- Keys do not have to match on SMLT core switches if 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. After 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 port values that differ from the global values.

Avaya Virtual Services Platform 9000 software uses the following LACP timers:

- fast periodic timer—200 to 20 000 milliseconds (ms); default 1000 ms
- slow periodic timer—10 000 to 30 000 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. Link Aggregation Control Protocol timers do not reset until the next time you restart LACP globally or on a port. This ensures consistency with peer switches.

After you enable LACP on a port, the timer values are set at the port level. You must toggle the LACP status after 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. You must partner passive mode ports with active mode ports for aggregation to occur.

Link aggregation scaling

The following table summarizes the number of aggregation groups supported on Virtual Services Platform 9000 modules. For more information about the latest scalability information, see the release notes included with your switch.

Table 1: Link aggregation scaling for the Virtual Services Platform 9000

Feature	Number
Aggregation groups <ul style="list-style-type: none">• 802.3ad aggregation groups (LAG)• multilink trunk groups	512
Maximum number of links (ports) for each MLT group	16

LACP configuration using ACLI

This section describes how to configure and manage link aggregation using the Avaya Command Line Interface (ACLI), including Link Aggregation Control Protocol (LACP), to increase the link speed and redundancy for higher availability.

MultiLink Trunking (MLT) with 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.

You can describe the LACP 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).

Before you begin

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

Important:

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

- After you make a timer change, restart the 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 action ensures consistency with peer switches.

- Virtual Services Platform 9000 does not support standby mode for LACP aggregation groups of larger than eight ports.

Configuring global LACP parameters

Configure LACP parameters globally. After you configure the LACP system priority globally, it applies to all LACP-enabled aggregators and ports. After you enable the LACP on an aggregator or a port, it uses the global system priority value.

The new global parameter configuration is in effect after the LACP is restarted globally or on each port.

About this task

Important:

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

Procedure

1. Log on to Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Change the system priority:

```
lACP system-priority <0-65535>
```

3. Configure additional LACP parameters as required:

```
lACP enable [aggr-wait-time <200-2000>] [fast-periodic-time
<200-20000>] [slow-periodic-time <10000-30000>] [smlt-sys-id
<0x00:0x00:0x00:0x00:0x00:0x00>] [timeout-scale <2-10>]
```

If you do not configure the optional parameters, the system uses the default values.

Example

```
VSP-9012:1(config)#lACP fast-periodic-time 2000
VSP-9012:1(config)#lACP enable
```

Variable definitions

Use the data in the following table to use the `lACP` command.

Variable	Value
aggr-wait-time <200–2000>	Configures the aggregation wait time (in milliseconds) globally. The default is 2000.
enable	Enables LACP globally. The default is disabled.
fast-periodic-time <200–20000>	Configures the fast periodic time (in milliseconds) globally. The default is 1000.
slow-periodic-time <10000–30000>	Configures the slow periodic time globally. The default is 30000.
smlt-sys-id <0x00:0x00:0x00:0x00:0x00:0x00>	Configures the LACP system ID globally. Enter a MAC address in the following format: 0x00:0x00:0x00:0x00:0x00:0x00.
system-priority <0-65535>	Configures the global LACP system priority. The higher the value, the lower the priority. The default is 32768.
timeout-scale <2-10>	Configures the timeout scale globally. The default value is 3.

Configuring LACP on a port

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

About this task

Important:

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

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

Procedure

1. Log on to Interface Configuration mode:

```
enable
```

```
configure terminal
interface GigabitEthernet {slot/port[-slot/port][,...]}
```

2. Change the LACP mode:

```
lacp mode <active|passive>
```

3. Change the port priority:

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

4. Optionally, change the system priority:

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

5. Configure aggregation for the port:

```
lacp agr-wait-time <200-2000>[aggregation enable]
```

If you do not configure the optional parameter, the system uses the default values.

6. Configure parameters for the partner device at the opposite end of the link:

```
lacp partner-key <0-65535>
```

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

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

```
lacp partner-state <0-255 | 0x0-0xff>
```

```
lacp partner-system-id 0x00:0x00:0x00:0x00:0x00:0x00
```

```
lacp partner-system-priority <0-65535>
```

7. Configure additional LACP parameters as required:

```
lacp enable [fast-periodic-time <200-20000>] [key <0-65535>] [slow-
periodic-time <10000-30000>] [timeout-time <long|short>] [timeout-
scale <2-10>]
```

If you do not configure the optional parameters, the system uses the default values.

Example

Configure LACP on ports 4/21 and 4/22:

```
VSP-9012:1(config)#interface gigabitethernet 4/21-4/22
VSP-9012:1(config-if)#lacp key 1 timeout short
VSP-9012:1(config-if)#lacp aggregation enable
VSP-9012:1(config-if)#lacp partner-key 0
VSP-9012:1(config-if)#lacp partner-port-priority 0
VSP-9012:1(config-if)#lacp enable
```

Variable definitions

Use the data in the following table to use the `lacp` command.

Variable	Value
<code>aggr-wait-time <200–2000></code>	Configures the aggregation wait time (in milliseconds) for this port. The default is 2000.
<code>aggregation enable</code>	Enables aggregation on the port, which makes it an aggregated link.
<code>enable</code>	Enables LACP for this port. The default is disabled.
<code>fast-periodic-time <200–20000></code>	Configures the fast periodic time (in milliseconds) for this port. The default is 1000 ms.
<code>key <0–65535></code>	Configures the aggregation key for this port.
<code>mode {active passive}</code>	Configures the LACP mode to be active or passive.
<code>partner-key <0–65535></code>	Configures the partner administrative key.
<code>partner-port <0–65535></code>	Configures the partner administrative port value.
<code>partner-port-priority <0–65535></code>	Configures the partner administrative port priority value.
<code>partner-state <0-255 0-255></code>	<p>Configures 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.</p> <p>For example, to set the two partner-state parameters</p> <ul style="list-style-type: none"> • Act = true • Agg = true <p>specify a value of 0x05 (bitmap = 00000101).</p>
<code>partner-system-id <0x00:0x00:0x00:0x00:0x00:0x00></code>	Configures the partner administrative system ID. Specify a MAC address in the format 0x00:0x00:0x00:0x00:0x00:0x00.
<code>partner-system-priority <0–65535></code>	Configures the partner administrative system priority value.
<code>priority <0–65535></code>	Configures the port priority. The default value is 32768. To set this option to the default value, use the <code>default</code> operator with the command.
<code>slow-periodic-time <10000-30000></code>	Configures the slow periodic time for this port. The default is 30000 ms. To set this option to the default value, use the <code>default</code> operator with the command.
<code>system-priority <0-65535></code>	Configures the system priority for this port. The default is 32768.
<code>timeout-scale <2-10></code>	<p>Configures a timeout scale for this port. The default value is 3.</p> <p>The LACP timeout is equal to the slow periodic time or fast periodic time multiplied by the timeout-scale, depending how you configure the timeout-time variable.</p>
<code>timeout-time {long short}</code>	Configures the timeout to either long or short.

Viewing LACP configuration information

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

Procedure

1. View the global configuration:

```
show lacp
```

2. View LACP administrative information for the actor device:

```
show lacp actor-admin interface [gigabitethernet]
```

OR

```
show lacp actor-admin interface gigabitethernet [vid <1-4084>]
[{{slot/port[-slot/port][, ...}}]
```

3. View LACP operational information for the actor device:

```
show lacp actor-oper interface [gigabitethernet]
```

OR

```
show lacp actor-oper interface gigabitethernet [vid <1-4084>]
[{{slot/port[-slot/port][, ...}}]
```

4. View LACP timer information:

```
show lacp extension interface [gigabitethernet]
```

5. View LACP interface configuration information

```
show lacp interface
```

OR

```
show lacp interface gigabitethernet [vid <1-4084>] [{{slot/port[-
slot/port][, ...}}]
```

OR

```
show lacp interface mlt [<64-6399>]
```

OR

```
show lacp interface mlt [id<1-512>]
```

6. View LACP administrative information for the partner device:

```
show lacp partner-admin interface [gigabitethernet]
```

OR

```
show lacp partner-admin interface gigabitethernet [vid <1-4084>]
[{{slot/port[-slot/port][, ...}}]
```

7. View LACP operational information for the partner device:

```
show lacp partner-oper interface [gigabitethernet]
```

OR

```
show lacp partner-oper interface gigabitethernet [vid <1-4084>]
[{slot/port[-slot/port][, ...]}]
```

Example

```
VSP-9012:1#show lacp
```

```
=====
                        Lacp Global Information
=====
      SystemId: 00:24:7f:a1:70:00
      SmltSystemId: 00:00:00:00:00:00
      Lacp: enable
      system-priority: 32768
        timeout-admin: 3
      fast-periodic-time-admin: 1000
      slow-periodic-time-admin: 30000
      aggr-wait-time-admin: 2000
        timeout-oper: 3
      fast-periodic-time-oper: 2000
      slow-periodic-time-oper: 30000
      aggr-wait-time-oper: 2000
```

In the following example output, `aggr` indicates the port has become part of an aggregation. `indi` indicates individual.

```
VSP-9012:1(config-if)#show lacp actor-admin interface gigabitethernet 4/21-4/22
```

```
=====
                        Actor Admin
=====
INDEX  SYS   SYS      KEY  PORT  PORT  STATE
      Prio ID
-----
4/21   32768 00:24:7f:9f:c0:00 1    0x114  32768 act short aggr
4/22   32768 00:24:7f:9f:c0:00 1    0x115  32768 act short aggr
```

```
VSP-9012:1(config-if)#show lacp partner-admin interface gigabitethernet 4/21-4/22
```

```
=====
                        Partner Admin
=====
INDEX  SYS   SYS      KEY  PORT  PORT  STATE
      Prio ID
-----
4/21   0     00:00:00:00:00:00 0    0x0    0    pas    long indi
4/22   0     00:00:00:00:00:00 0    0x0    0    pas    long indi
```

Variable definitions

Use the data in the following table to use the `show lacp` command.

Variable	Value
actor-admin interface gigabitethernet [vid <1-4084> [{slot/port[-slot/port][,...]}]]	Shows LACP actor administrative information for all interfaces or the specified interface. <ul style="list-style-type: none"> • <1-4084> is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • {slot/port[-slot/port][,...]} is the port or port list.
actor-oper interface gigabitethernet [vid <1-4084> [{slot/port[-slot/port][,...]}]]	Shows LACP actor operational information for all interfaces or the specified interface. <ul style="list-style-type: none"> • <1-4084> is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • {slot/port[-slot/port][,...]} is the port or port list.
extension interface gigabitethernet [vid <1-4084> [{slot/port[-slot/port][,...]}]]	Shows LACP timer information for all interfaces or the specified interface. <ul style="list-style-type: none"> • <1-4084> is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • {slot/port[-slot/port][,...]} is the port or port list.
interface [gigabitethernet [vid <1-4084> [{slot/port[-slot/port][,...]}]] [mlt <64-6399>] [mlt id <1-512>]]	Shows all LACP port configuration information for all interfaces or the interface you specify. <ul style="list-style-type: none"> • <1-4084> is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • {slot/port[-slot/port][,...]} is the port or port list. • <64-6399 is the interface index of the mlt. • <1-512 is the MLT ID.
partner-admin interface [gigabitethernet] [vid <1-4084> [{slot/port[-slot/port][,...]}]]	Shows LACP partner administrative information for all interfaces or the specified interface. <ul style="list-style-type: none"> • <1-4084> is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • {slot/port[-slot/port][,...]} is the port or port list.

Table continues...

Variable	Value
partner-oper interface [gigabitethernet] [vid <1-4084> [{slot/port[-slot/port]}[,...]]	Shows LACP partner operational information for all interfaces or the specified interface. <ul style="list-style-type: none"> • <1-4084> is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • {slot/port[-slot/port]}[,...] is the port or port list.

LACP configuration using EDM

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:

Virtual Services Platform 9000 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). Link Aggregation Control Protocol (LACP) can dynamically add or remove LAG ports, depending on their availability and states.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **LACP Global** tab.
4. To enable LACP globally, select the **Enable** check box.
5. Configure the remaining parameters as required.

Important:

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

6. Click **Apply**.

LACP Global field descriptions

Use the data in the following table to use the **LACP Global** tab.

Name	Description
Enable	Enables or disables LACP globally.
SystemPriority	Configures the system priority for all LACP enabled aggregators and ports. The default value is 32768.
FastPeriodicTime	Configures 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	Displays the operating value of the fast periodic timer on the port. The default value is 1000.
SlowPeriodicTime	Configures 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	Displays the operating value of the slow periodic timer on the port. The default value is 30000.
AggrWaitTime	Configures 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	Displays the operating value of the aggregate wait timer on the port. The default value is 2000.
TimeoutScale	Configures 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	Displays the operating value of the timeout scale on the port. The default value is 3.
SmltSysId	Configures the LACP system ID for SMLT. Use this optional parameter only for SMLT situations. You must configure the same LACP SMLT system ID on both

Table continues...

Name	Description
	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.

About this task

! **Important:**

The Virtual Services Platform 9000 does not support standby mode for aggregation groups of larger than eight ports.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **LACP** tab.
4. Double-click a field to change the value.
 You cannot edit grey-shaded fields in the table.
5. Click **Apply**.

LACP field descriptions

Use the data in the following table to use the **LACP** tab.

Name	Description
Index	The unique identifier the local system allocates to this aggregator. 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.
ActorSystemPriority	The two octet read-write value indicating the priority value associated with the actor 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

Table continues...

Name	Description
	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 permits 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, which 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, which 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, which 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, which 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

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand 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**.

LACP field descriptions

Use the data in the following table to use the **LACP** tab.

Name	Description
AdminEnable	Enables LACP status for the port. The default value is false.
OperEnable	Displays 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	Displays 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. The default value is 30000.
SlowPeriodicTimeOper	Displays 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	Displays 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	Displays the operating value of the timeout scale on the port. The default value is 3.

Table continues...

Name	Description
ActorSystemPriority	<p>Specifies the two octet read-write value indicating the priority value associated with the actor system ID. The range is 0–65535.</p> <p>The default value is 32768.</p>
ActorSystemID	<p>Displays the six octet read-write MAC address value used as a unique identifier for the system that contains this aggregator.</p> <p>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.</p>
ActorAdminKey	<p>Specifies the current administrative value of the key for the aggregator. The administrative key value can differ from the operational key value, which is a 16-bit read-write value in the range of 0–65535. This key needs to match the LAG key.</p>
ActorOperKey	<p>Displays the current operational value of the key for the aggregator. The administrative key value can differ from the operational key value, which is a 16-bit read-only value. The meaning of particular key values is of local significance.</p>
SelectedAggID	<p>Displays 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.</p>
AttachedAggID	<p>Displays 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.</p>

Table continues...

Name	Description
ActorPort	Displays 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	<p>Specifies the priority value assigned to this aggregation port. This 16-bit value is read-write in the range of 0–65535.</p> <p>The default value is 32768.</p>
ActorAdminState	<p>Specifies a string of eight bits, corresponding to the administrative values as transmitted by the actor in LACPDUs, by selecting check boxes. The values are</p> <ul style="list-style-type: none"> • the first bit corresponds to bit 0 of Actor_State (LACP_Activity) (the default value) • the second bit corresponds to bit 1 (LACP_Timeout) • the third bit corresponds to bit 2 (Aggregation) • the fourth bit corresponds to bit 3 (Synchronization) • the fifth bit corresponds to bit 4 (Collecting) • the sixth bit corresponds to bit 5 (Distributing) • the seventh bit corresponds to bit 6 (Defaulted) • the eighth bit corresponds to bit 7 (Expired) <p>These values allow administrative control over the values of LACP_Activity, LACP_Timeout, and aggregation. This attribute value is read-write.</p>
ActorOperState	<p>Displays a string of eight bits, corresponding to the current operational values of Actor_State as transmitted by the actor in LACPDUs. This attribute value is read-only.</p> <p>The default value is lacpActive.</p>
PartnerAdminSystemPriority	<p>Specifies the current administrative value of the port number for the protocol partner. It 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.</p> <p>The default value is 0.</p>
PartnerOperSystemPriority	Displays 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

Table continues...

Name	Description
	<p>PartnerAdminSystemPriority if there is no protocol partner.</p> <p>The default value is 0.</p>
PartnerAdminSystemID	<p>Specifies 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.</p> <p>The default value is 00:00:00:00:00:00.</p>
PartnerOperSystemID	<p>Displays a six octet read-only MAC address value that indicates representing the current value of the aggregation port 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.</p> <p>The default value is 00:00:00:00:00:00.</p>
PartnerAdminKey	<p>Specifies the current administrative value of the key for the protocol partner. It 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.</p> <p>The default value is 0.</p>
PartnerOperKey	<p>Displays the current operational value of the key for the aggregator current protocol partner. It is a 16-bit read-only value. If the aggregation is manually configured, this value is assigned by the local system.</p> <p>The default value is 0.</p>
PartnerAdminPort	<p>Specifies the current administrative value of the port number for the protocol partner. It 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.</p> <p>The default value is 0.</p>
PartnerOperPort	<p>Displays the operational port number assigned to this aggregation port by the aggregation port protocol partner. The value of this attribute can contain the</p>

Table continues...

Name	Description
	<p>manually configured value carried in AggPortPartnerAdminPort if there is no protocol partner. This 16-bit value is read-only.</p> <p>The default value is 0.</p>
PartnerAdminPortPriority	<p>Specifies the current administrative value of the port priority for the protocol partner. It 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.</p> <p>The default value is 0.</p>
PartnerOperPortPriority	<p>Displays 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.</p> <p>The default value is 0.</p>
PartnerAdminState	<p>Specifies a string of eight bits, corresponding to the current administrative value of Actor_State for the protocol partner, by selecting check boxes. This attribute value is read-write. The assigned value is used to achieve manually configured aggregation.</p> <p>The default value is none.</p>
PartnerOperState	<p>Displays a string of eight bits, corresponding to the current values of Actor_State in the most recently received LACPDU transmitted by the protocol partner. In the absence of an active protocol partner, this value can reflect the manually configured value PartnerAdminState. This attribute value is read-only.</p> <p>The default value is none.</p>

Chapter 4: MultiLink Trunking and Split MultiLink Trunking

This section provides the concepts and procedures you need to configure MultiLink Trunking (MLT) and Split MultiLink Trunking (SMLT).

Link aggregation overview

Link aggregation provides link level redundancy and increases load sharing. Use Link aggregation to bundle the ports into a port group, which is represented as one logical interface to the Media Access Control (MAC) layer.

The Avaya Virtual Services Platform 9000 supports the following types of link aggregation:

- MultiLink Trunking (MLT)—a statically configured link bundling method. MLT is not standards based, but it interoperates with static link methods of other vendors. If you aggregate links from multiple modules together, MLT can provide not only port redundancy, but also module redundancy.
- IEEE 802.3ad based link aggregation, through the Link Aggregation Control Protocol (LACP), dynamically aggregates links as they become available to a trunk group. Link Aggregation Control Protocol dynamically detects whether links can be aggregated into a link aggregation group (LAG) and does so after links become available. Link Aggregation Control Protocol also provides link integrity checking at Layer 2 for all links within the LAG.

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

The Virtual Services Platform 9000 software offers LACP functionality layered with MLT. This document uses the term MLT with LACP to refer to this functionality.

An Interswitch Trunk (IST) is a special form of aggregate link. ISTs use a link aggregation mechanism such as a Multilink Trunk (MLT) to create a single logical switch out of two physical devices. Creating one logical switch allows the devices to share resources, enhance fault tolerance, and create network redundancy.

Split MultiLink Trunking (SMLT) is an option that improves Layer 2 (bridged) resiliency by providing for the addition of a Virtual Services Platform 9000 failure redundancy with subsecond failover, on top of all standard MLT link failure protection and flexible bandwidth scaling functionality. Use SMLT to connect a device that supports some form of link aggregation, be it a switch or a server, to two

distinct separate SMLT endpoints or switches. These SMLT devices form a Switch Cluster and are referred to as an Interswitch Trunk (IST) Core Switch pair.

You can also use LACP on SMLT configurations. The Virtual Services Platform 9000 provides modifications to the LACP in SMLT configurations. LACP-capable devices can connect to an SMLT aggregation pair. Avaya recommends that you do not configure LACP on the IST MLT.

Virtual LACP (VLACP) is an Avaya modification that provides end-to-end failure detection. The VLACP is not a link aggregation protocol; VLACP implements link status control protocol at the port level. This mechanism periodically checks 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 an MLT. Avaya recommends that you do not configure VLACP on LACP-enabled ports. Virtual LACP does not operate properly with LACP. You can configure VLACP with an 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 aggregated bandwidth. Grouping multiple ports into a logical link provides a higher aggregate on a switch-to-switch or switch-to-server application.

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

MLT traffic distribution algorithm

You can use a multilink trunk to aggregate bandwidth between two switches. The Virtual Services Platform 9000 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 a flow always traverses the same link path.

The algorithms are the same traffic distribution algorithms used for the IEEE 802.3ad based link aggregation.

For more information about port and algorithm calculations, see [Viewing the MLT port calculated by the MLT hash algorithm](#) on page 53.

The Virtual Services Platform 9000 uses an enhanced traffic distribution algorithm, which ensures proper traffic distribution in all customer networks. The system determines MultiLink Trunking hashing for a module type at the ingress port, although the effect 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 the ingress port for the traffic, not the MLT configuration.

The hashing 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)]

For IPv6, the MLT hashing algorithm is:

64-bit key = SrcIP(LSB 64 bits) XOR SrcIP(MSB 64 bits) XOR DstIP(LSB 64 bits) XOR DstIP(MSB 64 bits)

For non IP/IPv6 traffic, the MLT hashing algorithm is:

64-bit key = [Dst Mac (LSB 32 bits), Src Mac (LSB 32 bits)]

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

Use the `hash-calc getmltindex` command to configure src-port and dst-port as optional parameters.

The hashing for IPv4 traffic between a given source and destination IP address is different for TCP/UDP packets and ICMP packets.

MultiLink trunking and autonegotiation interaction

To use MLT with the Virtual Services Platform 9000, you can have 10Gb and 1Gb ports running at different speeds. After you use MLT with LACP, LACP dynamically checks for proper speed on all port members. You do not need to have similar physical connection types. For example, you can mix a fiber port with a copper port. After you use autonegotiation with MLT and not LACP, you need to ensure that all ports run at the same speed.

Multicast flow distribution over MLT

MultiLink Trunking 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 after a large number of TV streams are employed.

With 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 more information about multicast flow distribution over MLT, see *Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000*, NN46250-504.

Multicast distribution algorithm

The multicast hashing algorithm is the same algorithm used for an IP packet. For more information, see [MLT traffic distribution algorithm](#) on page 30.

MLT and MLT with LACP configuration rules

Virtual Services Platform 9000 multilink trunks adhere to the following rules. Unless otherwise stated, these rules also apply to MLT with LACP.

- MLT is supported on 10Base-T, 100Base-TX, Gigabit Ethernet, and 10 Gigabit Ethernet module ports.
- Multilink trunk ports support mixed speed links, for example, one link can be 10Gb and another 1Gb. However, no weighting of traffic distribution occurs so if you mix links of different operational speeds, you can overload the lower speed link or under utilize a higher speed link.

This rule applies to multilink trunks only. MLT with LACP does not support different link speeds.

- 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. Use tagging so ports can belong to multiple STGs, as well as multiple VLANs.
- After 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. After you remove the port from MLT or after you delete the MLT, the ports are removed from the MLT STG and added into the default STG.
- MLT is compatible with 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.

Virtual Services Platform 9000 multilink trunks have the following general features and requirements:

- Supports up to 512 MLT groups with as many as 16 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 MSTP or RSTP enabled, ports in the same multilink trunk operate as follows:

- The designated port sends the 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.
- 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 Virtual Services Platform 9000 Link Aggregation Groups (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 must belong to the same set of VLANs.
- Link aggregation is compatible with MSTP, and RSTP.
- Assign all LAG ports to the same MSTP or RSTP groups.
- Ports in a LAG can exist on different modules.
- Support a maximum of 32 link aggregation groups.
- You can configure a LAG with up to 16 ports, but only a maximum of 8 can be active at a time.
- After you configure a multilink trunk 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 ports and port memberships to form a link aggregation group (LAG). Use Link Aggregation Control Protocol to gather one or more links to form a LAG, which a Media Access Control (MAC) client treats as a single link. Link Aggregation Control Protocol 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 or multiplexer, frame distribution, and frame collection functions.

The following illustration shows the major functions of IEEE 802.3ad defined as multiple link aggregation.




Figure 1: Link aggregation sublayer (according to IEEE 802.3ad)

The link aggregation sublayer comprises 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. Virtual Services Platform 9000 only implements the Marker Receiver function. For

more information about the frame distribution algorithm, see [MLT traffic distribution algorithm](#) on page 30.

- frame collection

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

- aggregator parser or 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 (PDU), 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 or 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 (LACPDU) from other frames. The blocks pass, passing the LACPDU to the appropriate sublayer entity and all other frames to the aggregator.

802.3ad link aggregation principles

Use link aggregation to group ports together to form a link group to another device. 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 if 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.
- 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. Because the frames are not ordered on a single link, this guarantees that frame ordering is maintained for all conversations.
- Conversations move among ports within an aggregation for load balancing and for maintaining availability if a link fails.
- The standard does not impose a particular distribution algorithm on the distributor. The Virtual Services Platform 9000 uses the MLT traffic distribution algorithm. For more information on the Virtual Services Platform 9000 traffic distribution algorithm, see [MLT traffic distribution algorithm](#) on page 30.
- Each port is assigned a unique, globally administered MAC address.

After 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

Split MultiLink Trunking (SMLT) is an option that improves Layer 2 and Layer 3 resiliency. The following sections discuss SMLT in more detail.

SMLT overview

Split MultiLink Trunking is an option that improves Layer 2 (bridged) resiliency by providing for the addition of Virtual Services Platform 9000 failure redundancy with subsecond failover, on top of all standard MLT link failure protection and flexible bandwidth scaling functionality. Use Split MultiLink Trunking to connect a 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 devices form a Switch Cluster and are referred to as an interswitch trunk (IST) Core Switch pair.

Switch Clusters are always formed as a pair, but you can combine pairs of clusters in either a square or full-mesh fashion to increase the size and port density of the Switch Cluster. If you configure SMLT in a Layer 3 or routed topology, the configuration is referenced as Routed SMLT (RSMLT). For information about Routed SMLT, see *Configuring IP Routing on Avaya Virtual Services Platform 9000*, NN46250-505.

You can form SMLT connections through single links from the Switch Cluster to the edge connection, standard MLTs, or MLTs with LACP. Optionally, SMLT links can have VLACP enabled as well. You can mix these various link connections. Within the same Switch Cluster, you can configure both SMLT and RSMLT to allow a mixture of both Layer 2 and Layer 3 VLANs. For more information about various SMLT configurations, see [SMLT topologies](#) on page 37.

Split MultiLink Trunking networks do not need to use RSTP or MSTP to enable loop-free triangle topologies because SMLT inherently avoids loops due to its superior enhanced link aggregation protocol. The loop-free link is accomplished by having two aggregation switches appear as a single

device to edge switches, which dual-home to the aggregation switches. The aggregation switches interconnect using an interswitch trunk, exchanging addressing and state information (permitting rapid fault detection and forwarding path modification). Split MultiLink Trunking is designed for Layer 2 network connectivity, but you can configure in Layer 3 networks by working with VRRP or RSMLT Layer 2 edge.

For more information about SMLT configuration, see *Switch Clustering using Split-MultiLink Trunking (SMLT) with VSP 9000, ERS 8600/8800, 8300, and 5000 Technical Configuration Guide*, NN48500-518. For more information about the technical configuration guide, go to the Avaya Web site: <http://support.avaya.com>.

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 failure
- elimination of single points of failure
- fast recovery in case of node failure
- transparent and interoperable solutions
- removal of MSTP and RSTP convergence issues

SMLT, MSTP, and RSTP

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 MSTP or RSTP for loop detection. A properly designed SMLT network inherently does not have logical loops.

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

If you use STP mode on a switch that is in an SMLT configuration with Virtual Services Platform 9000 systems, you can experience traffic loss for 30 seconds if you change the port membership of the MLT, even if the port is in a down state. The traffic loss is because the convergence time for STP is 30 seconds. Avaya recommends that you use MSTP or RSTP on all switches in SMLT configurations.

SMLT topologies

The following are the three generic topologies in which you can deploy SMLT, depending on the resiliency and redundancy required:

- a triangle topology
- a square topology
- a full-mesh topology

SMLT and interswitch trunking

The following illustration shows an SMLT configuration with a pair of Virtual Services Platform 9000 systems as aggregation systems (E and F) and four separate Virtual Services Platform 9000 systems as user access switches (A, B, C, and D).




Figure 2: SMLT aggregation switches and operations

You must connect SMLT aggregation switches through an interswitch trunk. For example, user access switches B and C connect to the aggregation systems through multilink trunks split between the two aggregation systems. As shown in [Figure 2: SMLT aggregation switches and operations](#) on

page 37, the implementation of SMLT only requires two SMLT-capable aggregation systems. You must connect these Virtual Services Platform 9000 systems through an interswitch trunk.

Aggregation switches use the interswitch trunk to do the following:

- 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.

If the SMLT interfaces from one of the IST peers are down, all MAC entries and IP Address Resolution Protocol (ARP) entries learned from SMLT interfaces are unchanged. This behavior is different than on Avaya Ethernet Routing Switch 8600, where all corresponding MAC and ARP entries are converted to be learned from the IST interface.

Avaya recommends that an interswitch multilink trunk contain at least two physical ports. You cannot enable the CP-Limit on all physical ports that are members of an interswitch trunk multilink trunk.

Important:

Avaya does not allow you to enable the CP-Limit on the IST MLT in the VSP 9000.

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. You can disable one of the multilink trunk ports (risers) in such a condition; traffic continues to flow uninterrupted through the remaining SMLT ports.

For more information about CP-Limit, see *Administering Avaya Virtual Services Platform 9000*, NN46250-600.

Other SMLT aggregation switch connections

The example shown in [Figure 2: SMLT aggregation switches and operations](#) on page 37 includes end stations that connect to each of the Virtual Services Platform 9000 systems. 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 a 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. The packet is routed correctly 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 an active SMLT group in order to prevent network loops.
- After a packet is received, the system performs a look-up on the forwarding database. If an entry exists, and if the entry was learned locally from the SMLT or through the IST as a remote SMLT, it is forwarded to the local port (the packet must not be sent to the IST 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 Virtual Services Platform 9000 obeys the MLT traffic distribution algorithm.

SMLT traffic flow examples

The following traffic flow examples are based on [Figure 2: SMLT aggregation switches and operations](#) on page 37.

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. The link 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 passes through one active VRRP Master for each IP subnet. The traffic 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. The entry for 00:E0:7B:B3:04:00 is shown on node C as an entry 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 directly learned through MLT-1 because SMLT Remote is false.

The following illustration shows the network topology.




Figure 3: Network topology for traffic flow example

After a packet arrives at node C destined for 00:E0:7B:B3:04:00, if the SMLT Remote status is true, the Virtual Services Platform 9000 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.

SMLT and LACP support

Virtual Services Platform 9000 fully supports IEEE 802.3ad LACP on MLTs and on a pair of SMLT systems.

With LACP the Virtual Services Platform 9000 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 the following:

- MLT peers and SMLT client devices can be both network switches and a type of server or 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 MLT IDs assigned to SMLT client. SMLT switches can detect if MLT 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.
- The 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 CP module failures. In the case of a CP module failure in a system with only one CP module, 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 for failure recovery for the network sends the data along a different network path.
- Link Aggregation Control Protocol (LACP) is not supported on Split Multilink Trunking (SMLT) Shortest Path Bridging MAC (SPBM) Network Node Interfaces (NNIs) port type because it can cause traffic to be silently dropped.

SMLT and IP routing

This section describes SMLT and IP routing interactions.

- [SMLT and the Virtual Router Redundancy Protocol](#) on page 41
- [SMLT and VRRP BackupMaster](#) on page 41

SMLT and the Virtual Router Redundancy Protocol

Use Virtual Router Redundancy Protocol (VRRP) to 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.

* Note:

Avaya does not support a VRRP virtual IP address to be the same as the local physical address of the device.

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. The system directly routes all traffic to the destined subnetwork and not through Layer 2 switches to the VRRP master. This avoids a potential limitation in the available IST bandwidth.

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

If you use 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 a 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 avoids traffic flow across the IST. This provides true load-sharing abilities.

The BackupMaster feature provides an additional benefit. Virtual Router Redundancy Protocol normally sends a hello packet every second. After three hello packets are not received, all switches automatically revert to master mode. This results in a 3-second outage. If you use VRRP in an SMLT environment, and a link goes down, the system automatically forwards traffic to the remaining ports that you configure as SMLT VRRP BackupMaster. Because both systems process traffic, the node immediately recognizes the VRRP state change, so there is faster failure recovery (less than 1 second).

Follow these guidelines if 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 one hundred and fifty 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 Virtual Services Platform 9000 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

You can use Simple Loop Prevention Protocol (SLPP) to prevent loops in an SMLT network. SLPP focuses on SMLT networks but works with other configurations. Avaya recommends that you always use SLPP in an SMLT environment.

For square and full-mesh configurations that use a routed core, Avaya recommends that you create a separate core VLAN. Enable SLPP on the core VLAN and the square or full mesh links between switch clusters. This configuration detects loops created in the core and loops at the edge do not affect core ports. If you use RSMLT between the switch clusters, enable SLPP on the RSMLT VLAN. Because you enable SLPP only on one or two VLANs in the core, changing the RX threshold values will not be necessary.

The Virtual Services Platform 9000 does not support SLPP in an SMLT that uses LACP.

The SLPP design is to shutdown the port where the SLPP packets originate or to shutdown the IST peer switch port, after the counter reaches the threshold. A loop can still occur after ports are shutdown. SLPP can shutdown all SMLT ports on a triangle SMLT topology, which results in isolating the edge switch.

For more information about SLPP fundamentals and configuration, see *Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000*, NN46250-500.

MLT and SMLT configuration considerations

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

MLT with LACP configuration considerations

After 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 to 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 MLT and assign to it the same key as in step 1.

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

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

After a multilink trunk is configured with LACP, you cannot add or delete ports manually without first disabling LACP. You can add or remove VLANs to an MLT without manually disabling LACP. When you add or remove VLANs from an LACP-enabled MLT, follow the guidelines in [Configuring LACP on an MLT](#) on page 49 or [Adding a multilink interswitch or LACP trunk](#) on page 56.

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

You can configure Split MultiLink Trunking (SMLT) with MLT, or with MLT with LACP. Follow these guidelines while you configure SMLT with LACP:

- If you configure LACP for SMLT, you must configure the same LACP smlt-sys-id on both sides.

After you configure the LACP system ID for SMLT, configure the same LACP smlt-sys-id on both aggregation switches to avoid loss of data. Avaya recommends that you configure the LACP smlt-sys-id to be the base MAC address of one of the aggregate switches, and that you include the MLT-ID. Configure the same system ID 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 you need to modify the LACP configuration, for example, change the LACP key, on a square or full-mesh core network, you must perform the following tasks:
 1. Shutdown all LACP ports.
 2. Make the configuration change.
 3. Bring all LACP ports back up.

If you only disable LACP without shutting down the LACP ports, you can create a looping scenario.

- If an SMLT aggregation switch has LACP enabled on some of its multilink trunks, do not change the LACP system priority.
- After 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 ISTs to avoid unnecessary processing. Use VLACP if an optical network between the SMLT core switches requires a failure detection mechanism.
- Enabling multicast MLT distribution on a square-mesh SMLT can cause loss of streams going through the diagonal links. This results in traffic distributed improperly. Avaya recommends that you disable the diagonal links.

Using the LACP `smlt-sys-id` enables you to use a 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 always uses the same LACP key regardless of the state of the SMLT Core Aggregation Switch neighbor (or the IST link). Therefore no change in LAGs must occur on the attached device 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 MLT-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 MLT-ID. The administrator must ensure that the same value for the system ID is configured on both of the SMLT Core Aggregation Switches.

The Virtual Services Platform 9000 software does not support Simple Loop Prevention Protocol (SLPP) in an LACP-SMLT environment.

You can configure the LACP `smlt-sys-id` used by SMLT core aggregation switches. After you set the LACP system ID for SMLT, configure the same LACP `smlt-sys-id` on both aggregation switches to avoid the loss of data.

The LACP System ID is the base MAC address of the switch, which is carried in Link Aggregation Control Protocol Data Units (LACPDU). If 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 you enable the links for aggregation using the same key, LACP can dynamically aggregate them into an MLT LAG.

If 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 restarts, 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, after the restarted switch comes back on line, the same actions occur, thus disrupting traffic twice.

The solution to this problem is to statically configure the same LACP smlt-sys-id MAC address on both aggregation switches.

*** Note:**

The SMLT ID is always the same as the MLT ID. For instance, both sides can have an MLT 10, but once SMLT is enabled on both sides it will function as an SMLT. Until SMLT is enabled on both peers however, it will function as a normal MLT.

MLT with LACP and Spanning Tree configuration considerations

Only the physical link state or its LACP peer status affects LACP module operation. After a link is enabled or disabled, an LACP module is notified. The MSTP or RSTP forwarding state does not affect LACP module operation. LACPDUs can be sent if the port is in an MSTP or RSTP 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 that the configuration is created.

! Important:

Link Aggregation Control Protocol, as defined by IEEE, is a protocol that exists between two bridge endpoints; therefore, the LACP PDUs are terminated at the next Server Provider (SP) interface.

LACP parameters configuration considerations

Link aggregation scaling

The following table summarizes the number of aggregation groups supported on Virtual Services Platform 9000 modules. For more information about the latest scalability information, see the release notes included with your switch.

Table 2: Link aggregation scaling for the Virtual Services Platform 9000

Feature	Number
Aggregation groups <ul style="list-style-type: none"> • 802.3ad aggregation groups (LAG) • multilink trunk groups 	512
Maximum number of links (ports) for each MLT group	16

SMLT with NNI ports configuration considerations

If you want to modify SMLTs that contain NNI ports, Avaya recommends you do the modification during maintenance windows. Otherwise, if you create or delete SMLTs that contain NNI ports on Virtual Services Platform 9000 running MSTP, IS-IS adjacencies that connect to those ports can bounce even if the SMLT is not used.

MLT and SMLT link aggregation configuration using the ACLI

This section describes how to configure and manage link aggregation using the ACLI, including MultiLink Trunking (MLT), to increase the link speed and redundancy for higher availability.

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

For more advanced configuration information, see *Switch Clustering using Split Multi-Link Trunking (SMLT) with VSP 9000, ERS 8600/8800, 8300, and 5000 Technical Configuration Guide* (NN48500-518). For more information about the technical configuration guide, go to the Avaya Web site: <http://support.avaya.com>.

Configuring an MLT

Perform this procedure to create and configure an MLT.

Procedure

1. Log on to Global Configuration mode:

```
enable
configure terminal
```

2. Create an MLT:

```
mlt <1-512>
```

3. Add a VLAN to the MLT:

```
mlt <1-512> vlan <1-4084>
```

4. Optionally, change the name of the MLT:

```
mlt <1-512> name WORD<0-20>
```

5. Enable trunking on the MLT:

```
mlt <1-512> encapsulation dot1q
```

6. Enable the MLT:

```
mlt <1-512> enable
```

7. Display the MLT configuration:

```
show mlt <1-512>
```

Example

Create MLT 1:

```
VSP-9012:1(config)#mlt 1
```

Add VLAN 10 to the MLT:

```
VSP-9012:1(config)#mlt 1 vlan 19
```

Enable the MLT:

```
VSP-9012:1(config)#mlt 1 enable
```

Display the MLT configuration:

```
VSP-9012:1(config)#show mlt 1
```

```
=====
                                Mlt Info
=====
```

MLTID	IFINDEX	NAME	PORT TYPE	MLT ADMIN	MLT CURRENT	PORT MEMBERS	VLAN IDS
1	6144	SMLT-1	trunk	smlt	norm	4/2,4/14,4/26,4/38	2 3 4 6 19
MLTID	IFINDEX	DESIGNATED PORTS	LACP ADMIN	LACP OPER			
1	6144	4/2	disable	down			
MLTID	NAME	WHERE CREATED	LOCAL PORT MEMBERS	REMOTE PORT MEMBERS	WHICH PORTS PROGRAMMED IN DATA PATH		
1	SMLT-1	LOCAL	4/2,4/14,4/26,4/38		LOCAL		
MLTID	IFINDEX	ENCAP DOT1Q	LOSSLESS				
1	6144	enable	disable				
MLTID	IFINDEX	PORT	PACKET-RATE	SHUTDOWN			
1	6144	4/2	8000	false			
1	6144	4/14	8000	false			
1	6144	4/26	8000	false			
1	6144	4/38	8000	false			

Variable definitions

Use the data in the following table to use the `mlt` command.

Variable	Value
enable	Creates and enables a new MLT.
encapsulation dot1q	Enables trunking on the MLT.
member {slot/port[-slot/port][,...]}	Adds ports to the MLT.
<1-512>	Specifies the MLT ID.
name WORD<0-20>	Configures the name for the MLT.
vlan <1-4084>	Specifies a VLAN ID to add to the MLT.

Adding ports to an MLT LAG

Perform this procedure to add ports to an MLT LAG.

Procedure

1. Log on to Global Configuration mode:

```
enable
configure terminal
```

2. Add ports to an MLT LAG:

```
mlt <1-512> member {slot/port[-slot/port][,...]}
```

Example

Add ports 4/26 to 4/28 to MLT 1:

```
VSP-9012:1(config)#mlt 1 member 4/26-4/28
```

Variable definitions

Use the data in the following table to use the `mlt` command.

Variable	Value
enable	Creates and enables a new MLT.
encapsulation dot1q	Enables trunking on the MLT.
member {slot/port[-slot/port][,...]}	Adds ports to the MLT.
<1-512>	Specifies the MLT ID.
name <i>WORD</i> <0-20>	Configures the name for the MLT.
vlan <1-4084>	Specifies a VLAN ID to add to the MLT.

Removing ports from an MLT LAG

Remove ports from an MLT LAG.

About this task

Important:

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

Procedure

1. Log on to Interface Configuration mode:

```
enable
```

```

configure terminal
interface GigabitEthernet {slot/port[-slot/port][,...]}

```

2. Disable the ports:

```

shutdown port {slot/port[-slot/port][,...]}

```

3. Remove ports from an MLT LAG:

```

no mlt <1-512> member {slot/port[-slot/port][,...]}

```

Example

Remove ports 4/27 and 4/28 from MLT 1:

```

VSP-9012:1(config)#interface mlt 1
VSP-9012:1(config-mlt)#shutdown port 4/27-4/28
VSP-9012:1(config-mlt)#no mlt 1 member 4/27-4/28

```

Variable definitions

Use the data in the following table to use the `mlt` command.

Variable	Value
<1-512>	Specifies which MLT to add the ports to.
{slot/port[-slot/port][,...]}	Specifies the ports to add to the MLT.

Configuring LACP on an MLT

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

About this task**! 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.

When you add a VLAN to a dynamic MLT, only the active ports of the MLT are added as port members of the VLAN. Ports configured with the same aggregation key, but not active, are not added to the VLAN. If these inactive ports become active later, the system does not automatically add them to the VLAN port member list.

You must add all inactive ports to the VLAN. If you do not add the inactive ports to the VLAN, when they become active later, hashing can result in choosing a newly active port for traffic forwarding. Because the port is not a port member of the VLAN, traffic will be dropped. Avaya recommends that when you add the VLAN to the MLT, you also add the inactive aggregation ports to the VLAN. You may have to disable LACP on the inactive ports before you can add them to the VLAN. Because the ports are inactive, disabling LACP does not cause a traffic interruption.

Similarly when you remove a VLAN from a dynamic MLT, all active ports of the MLT are removed from the VLAN port member list but the inactive members are not removed. You must remove the inactive aggregation members from the VLAN.

If you later configure a port for the same aggregation, you must add this port to all VLANs that are members of the MLT.

Procedure

1. Log on to the MLT Interface Configuration mode:

```
enable
configure terminal
interface mlt <1-512>
```

2. Configure LACP on an MLT:

```
lacp enable [key <0-512>] [system-priority <0-65535>]
```

Example

Enable LACP and configure the LACP key on MLT 3:

```
VSP-9012:1(config)#interface mlt 3
VSP-9012:1(config-mlt)#lacp enable key 1281
```

Variable definitions

Use the data in the following table to use the `lacp` command.

Variable	Value
enable	Enables LACP on the MLT interface.
key <0-512>	Configures the LACP aggregator key for a specific MLT. <ul style="list-style-type: none"> • 0-512 is the LACP actor admin key.
system-priority <0-65535>	Configures the LACP system priority for a specific MLT. <ul style="list-style-type: none"> • 0-65535 is the system priority.

Creating an SMLT from an existing MLT

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

* Note:

If you want to modify SMLTs that contain NNI ports, Avaya recommends you do the modification during maintenance windows. Otherwise, if you create or delete SMLTs that contain NNI ports

on Virtual Services Platform 9000 running MSTP, IS-IS adjacencies that connect to those ports can bounce even if the SMLT is not used.

Before you begin

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

Procedure

1. Log on to the MLT Interface Configuration mode:

```
enable
configure terminal
interface mlt <1-512>
```

2. Create an SMLT from an existing MLT:

```
smlt
```

Example

Create an SMLT on MLT 1:

```
VSP-9012:1(config)#interface mlt 1
VSP-9012:1(config-mlt)#smlt
```

Creating an interswitch trunk

Create an interswitch trunk from an existing MLT.

About this task

Important:

Do not enable loop-detect on IST ports.

If you attempt to enable loop-detect on an existing IST port, the system prevents you from doing so. However, if you have a port with loop-detect already enabled, and you add that port to an IST, the system does not prevent you from doing so, causing potential system errors.

Procedure

1. Log on to the MLT Interface Configuration mode:

```
enable
configure terminal
interface mlt <1-512>
```

2. Create an interswitch trunk:

```
ist peer-ip <A.B.C.D> vlan <1-4084>
```

Example

Create an interswitch trunk on MLT 2:

```
VSP-9012:1(config)#interface mlt 2
VSP-9012:1(config-mlt)#ist peer-ip 47.17.11.35 vlan 5
```

Variable definitions

Use the data in the following table to use the `ist peer-ip` command.

Variable	Value
<A.B.C.D>	Specifies the peer IP address—the IP address of the IST VLAN on the other aggregation switch.
<1–4084>	Specifies the VLAN ID for this IST.

Enabling an interswitch trunk

Perform this procedure to enable an interswitch trunk.

Procedure

1. Log on to the MLT Interface Configuration mode:

```
enable
configure terminal
interface mlt <1-512>
```

2. Enable an interswitch trunk:

```
ist enable
```

Example

Enable an interswitch trunk on MLT 2:

```
VSP-9012:1(config)#interface mlt 2
VSP-9012:1(config-mlt)#ist enable
```

Variable definitions

Use the data in the following table to use the `ist` command.

Variable	Value
enable	Enables IST on this MLT interface. Use the no form of this command to disable IST on this MLT interface: <code>no ist enable</code> .

Deleting an interswitch trunk

Delete an interswitch trunk.

Before you begin

! Important:

Before you can delete the interswitch trunk, you must disable the interswitch trunk.

* Note:

The `no ist enable` command only dynamically disables the IST. The IST will become enabled again the next time the chassis reboots. Use the following procedure to permanently delete an IST.

Procedure

1. Log on to the MLT Interface Configuration mode:

```
enable
configure terminal
interface mlt <1-512>
```

2. Disable an interswitch trunk:

```
no ist enable
```

3. Delete an interswitch trunk:

```
no ist peer-ip
```

Example

From Global Configuration mode, log on to the MLT interface:

```
VSP-9012:1(config)# interface mlt 2
```

Disable the IST:

```
VSP-9012:1(config)# no ist enable
```

Delete the IST:

```
VSP-9012:1(config)# no ist peer-ip
```

Viewing the MLT port calculated by the MLT hash algorithm

View the MLT port calculated by the MLT hash algorithm to determine through which MLT port a packet will exit the system.

Procedure

View the MLT port calculated by the MLT hash algorithm:

```
hash-calc getmltindex traffic-type <ipv4|ipv6|non-ip> dest-val
WORD<1-1536> src-val WORD<1-1536> mltID <1-512> [src-port <0-65535>]
[dest-port <0-65535>]
```

The source and destination addresses cannot have the same value.

Example

View the MLT port calculated by the MLT hash algorithm:

```
VSP-9012:1#hash-cal getmltindex traffic-type ipv4 dest-val 47.17.11.30
src-val 47.17.10.31 mltid 1
```

Traffic will egress port 4/26 for this flow for MltId 1

Variable definitions

Use the data in the following table to use the **hash-calc getmltindex** command.

Variable	Value
dest-port <0-65535>	Specifies the destination port.
dest-val WORD<1-1536>	Specifies the destination address. The source and destination addresses cannot have the same value.
mltID <1-512>	Specifies the MLT ID. The value ranges from 1–512.
src-port <0-65535>	Specifies the source port.
src-val WORD<1-1536>	Specifies the source address. The source and destination addresses cannot have the same value.
traffic-type {ipv4 ipv6 non-ip}	Specifies the type of traffic, either ipv4, ipv6, or non-ip.

Viewing information about collision errors

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

Procedure

View information about collision errors:

```
show mlt error collision [<1-512>]
```

Example

```
VSP-9012:1(config-mlt)#show mlt error collision 4
```

```
=====
                                Mlt Collision Error
=====
MLT  -----COLLISIONS-----
ID   SINGLE  MULTIPLE LATE  EXCESSIVE
-----
4    0        0        0        0
=====
```

Variable definitions

Use the data in the following table to use the `show mlt error collision` command.

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

Viewing information about Ethernet errors

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

About this task

! Important:

The IMAC columns refer to internal MAC address errors.

Procedure

View information about Ethernet errors:

```
show mlt error main [<1-512>]
```

Example

```
VSP-9012:1(config-mlt)#show mlt error main 4
```

```

=====
                                Mlt Ethernet Error
=====
MLT  ALIGN  FCS    IMAC   IMAC   CARRIER  FRAMES  SQETEST  DEFER
ID   ERROR  ERROR  TRNSMIT RECEIVE SENSE    TOOLONG ERROR   TRNSMSS
-----
4    0      0      0      0      0        0       0        0

```

Variable definitions

Use the data in the following table to use the `show mlt error main` command.

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

Viewing all ports configured for SMLT

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

Procedure

View all ports configured for Single Port SMLT:

```
show smlt mlt
```

Example

```
VSP-9012:1(config-mlt)#show smlt mlt
```

```
=====
                                     Mlt SMLT Info
=====
MLT   ADMIN   CURRENT
ID    TYPE     TYPE
-----
1     smlt     smlt
4     smlt     smlt
```

Variable definitions

Use the data in the following table to use the `show smlt` command.

Variable	Value
mlt	Displays SMLT information for the MLT interface.

MLT and SMLT link aggregation configuration using EDM

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

For more advanced configuration information, see *Switch Clustering using Split Multi-Link Trunking (SMLT) with VSP 9000, ERS 8600/8800, 8300, and 5000 Technical Configuration Guide* (NN48500-518). For more information about the technical configuration guide, go to the Avaya Web site: <http://support.avaya.com>.

Adding a multilink, interswitch, or LACP trunk

Perform this procedure to add a multilink, interswitch, or LACP trunk.

*** Note:**

If you want to modify SMLTs that contain NNI ports, Avaya recommends that you do the modification during maintenance windows. Otherwise, if you create or delete SMLTs that contain NNI ports on Virtual Services Platform 9000 running MSTP, IS-IS adjacencies that connect to those ports can bounce even if the SMLT is not used.

About this task

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 autonegotiation is used. This requirement is not enforced by the software.

When you add a VLAN to a dynamic MLT, only the active ports of the MLT are added as port members of the VLAN. Ports configured with the same aggregation key, but not active, are not added to the VLAN. If these inactive ports become active later, the system does not automatically add them to the VLAN port member list.

You must add all inactive ports to the VLAN. If you do not add the inactive ports to the VLAN, when they become active later, hashing can result in choosing a newly active port for traffic forwarding. Because the port is not a port member of the VLAN, traffic will be dropped. Avaya recommends that when you add the VLAN to the MLT, you also add the inactive aggregation ports to the VLAN. You may have to disable LACP on the inactive ports before you can add them to the VLAN. Because the ports are inactive, disabling LACP does not cause a traffic interruption.

Similarly when you remove a VLAN from a dynamic MLT, all active ports of the MLT are removed from the VLAN port member list but the inactive members are not removed. You must remove the inactive aggregation members from the VLAN.

If you later configure a port for the same aggregation, you must add this port to all VLANs that are members of the MLT.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. 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 **access** or **trunk**.
7. In the **Name** box, type a name for the MLT, or accept the default.
8. In the **PortMembers** box, click the (...) button.
9. In the **Port Editor: PortMembers** dialog box, select the desired ports.
10. Click **Ok**
11. In the **VlanIdList** box, click the (...) button,
12. In the **VlanIdList** dialog box, select the desired VLANs.
13. Click **Ok**.
14. In the **MltType** section, select the MLT type.
15. In the **Aggregatable** box, select **enable** or **disable**.
16. Click **Insert**.

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

MultiLink/LACP Trunks field descriptions

Use the data in the following table to use the **MultiLink/LACP Trunks** tab

Name	Description
Id	Specifies a unique value for this MLT.
PortType	Specifies the value to access or trunk port. If the aggregatable field is set to enable, this field is read-only. The default value is access.
Name	Configures the name given to the MLT.
PortMembers	Assigns ports 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 Virtual Services Platform 9000 modules, up to eight same-type ports can belong to a single MLT. If the aggregatable field is set to enable, this field becomes read-only. ! 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 autonegotiation is enabled.
VlanIdList	Indicates to which ports the VLANs belong. If the aggregation field is set to enable, this field is read-only.
MltType	Specifies the type of MLT <ul style="list-style-type: none"> • normalMLT (default) • istMLT • splitMLT
RunningType	Specifies the MLT running type.
IfIndex	Specifies the interface of the trunk
ClearLinkAggregate	Clears the link aggregate, disabling and re-enabling the trunk.
DesignatedPort	Specifies the designated port of this trunk.
Aggregatable	Enables or disables link aggregation.

Table continues...

Name	Description
	The default value is disable.
AggOperState	Specifies the aggregation state of the trunk.
AggTimeOfLastOperChange	Specifies the time value since the interface entered its current operational state.
PeerPortMembers	Specifies the peer ports connected to the local ports of this trunk.
CpLimitPktRate	Limits control packets to QoS levels 6 and 7. Configures the limit in packets per second.
CpLimitShutDownEnable	Indicates if the MLT ports shutdown after the value in CpLimitPktRate is exceeded.
EntryOwner	Defines the owner of the MLT.
DatapathProgrammingState	Defines the datapath programming state of the MLT.

Adding ports to an MLT

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

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. In the **PortMembers** column, double-click the field associated with the MLT to which you want to add ports to.
5. Select the port numbers to add, or click **All** to add all ports to the MLT.
Up to 16 same-type ports can belong to a single MLT.
6. Click **Ok**.
7. Click **Apply**.

Viewing trunks

Perform this procedure to view the MLT-based SMLT configuration.

Procedure

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

SMLT Info field descriptions

Use the data in the following table to use the **SMLT Info** tab.

Name	Description
Id	Shows the MLT ID for this SMLT Read-only.
MltType	Shows the MLT type of this trunk.
RunningType	Shows the SMLT running type.

Configuring an IST MLT

Configure an IST MLT to add an IST to MLT_LACP.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. Click **Insert**.
5. In the **PortType** box, select **trunk**.
6. In the **PortMembers** box for the IST MLT, click the (...) button.
7. Click the ports you want to include in the IST MLT.
8. Click **Ok**.
9. In the **MltType** box, select **istMLT**.

Important:

Do not enable loop-detect on IST ports.

If you attempt to enable loop-detect on an existing IST port, the system prevents you from doing so. However, if you have a port with loop-detect already enabled, and you add that port to an IST, the system does not prevent you from doing so, causing potential system errors.

10. Configure the remaining parameters as required.
11. Click **Insert**.
12. In the navigation tree, expand the following folders: **Configuration > VLAN**.
13. Click **MLT/LACP**.
14. Click the **MultiLink/LACP Trunks** tab.
15. In the **MltType** field, select the MLT that you just created.
16. Click the **istMlt** button.

17. Type the peer IP address and the VLAN ID in the box.
18. Click **enable**.
19. Click **Apply**.
20. In the Device Physical View tab, select all IST MLT ports.
21. In the navigation tree, open the following folders: **Configuration > Edit > Port**
22. Click **General**.
23. Click the **CP Limit** tab.
24. In the **CpLimitShutDownEnable** option for each port, select **false**.
25. Click **Apply**.

Editing an IST

Edit an IST to make changes to an IST MLT.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. Select the IST.
5. Click **IstMlt**.
6. In the **PeerIp** box, type the peer IP address.
7. In the **VlanId** box, enter a VLAN ID.
8. In the **SessionEnable** box, select either **enable** or **disable**.
9. Click **Apply**.

IstMlt field descriptions

Use the data in the following table to help you configure the **IstMlt** tab.

Name	Description
PeerIp	Specifies the peer IP address, which is the IP address of the IST VLAN on the other aggregation switch.
VlanId	Configures an IST VLAN ID number.
SessionEnable	Enables and disables IST functionality. The default value is disable.
SessionStatus	Displays the status of the IST session.

MLT configuration examples

This section contains configuration examples for configuring MultiLink Trunking (MLT) and MLT with Link Aggregation Control Protocol (LACP) using Avaya Command Line Interface (ACLI).

MultiLink Trunking

This configuration example shows you how to create a multilink trunk and a Virtual Local Area Network (VLAN) between two Virtual Services Platform 9000 systems.

The following illustration shows you an MLT within a VLAN used to carry user traffic.




Figure 4: MLT within a VLAN

VSP 9000 A and B configuration

Configure MLT 10 on VLAN 100 on each VSP device.

```
#
# MLT CONFIGURATION
#
mlt 10 enable
interface mlt 10
exit
#
# VLAN CONFIGURATION
#
vlan create 100 name "VLAN-MLT-10" type port-mstprstp 1 color 1
vlan mlt 100 1
vlan mlt 100 10
vlan members 100 4/1-4/2,5/1-5/2 portmember
```

MultiLink Trunking with Link Aggregation Control Protocol

This configuration example shows you how to configure and enable a multilink trunk with LACP.

You must configure all aggregatable ports to use the same aggregator key so they can form an MLT.

The following illustration shows you an MLT created with LACP within a VLAN used to carry user traffic.




Figure 5: MLT within a VLAN

The following procedures show you how to configure both Virtual Services Platform 9000 systems.

VSP 9000 A and B configuration

Configure the two devices to connect with MLT with LACP.

```
#
# MLT CONFIGURATION
#
mlt 10 enable
interface mlt 10
lACP enable key 10
exit
#
# VLAN CONFIGURATION
#
vlan create 100 name "VLAN-MLT-10" type port-mstprstp 1 color 1
vlan mlt 100 1
vlan mlt 100 10
vlan members 100 4/1-4/2,5/1-5/2 portmember

#
# PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 4/1
lACP key 10 aggregation enable
lACP enable
exit
interface GigabitEthernet 4/2
lACP key 10 aggregation enable
lACP enable
exit
interface GigabitEthernet 5/1
lACP key 10 aggregation enable
lACP enable
exit
interface GigabitEthernet 5/2
lACP key 10 aggregation enable
lACP enable
exit
```

MLT network topology and configuration reference

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

The following illustration shows two multilink trunks (MLT1 and MLT2) connecting VSP 9000 C to switches VSP 9000 A and VSP 9000 B.




Figure 6: Switch-to-switch multilink trunks configuration

In this example, you can configure each trunk with multiple switch ports to increase bandwidth and redundancy. If traffic between switch-to-switch connections approaches single port bandwidth limitations, you can create a 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

In this example, File server 1 uses dual MAC addresses, with one MAC address for each Network Interface Card (NIC). No multilink trunk is configured on File server 1. File server 2 is a single MAC server (with a four-port NIC) configured as multilink trunk configuration MLT1.

The following illustration shows a typical switch-to-server multilink trunk configuration.




Figure 7: Switch-to-server multilink trunk configuration

In this example, one port on File server 1 is blocked and unused and File server 2 benefits from aggregated bandwidth on multilink trunk T1.

Example 3: Client/Server MLT

In this example, both servers are connected directly to VSP 9000 A. File server 2 is connected through multilink trunk MLT 1. The switch-to-switch connections are through MLT 2, MLT 3, and MLT 4. Clients access data from the servers (File server 1 and File server 2) and receive maximized bandwidth through MLT 1, MLT 2, MLT 3, and MLT 4. On the Avaya Virtual Services Platform 9000 systems, trunk members (the ports that comprise each multilink trunk) use different modules, to provide module redundancy.

The following illustration shows how you can use multilink trunks in a client/server configuration.




Figure 8: Client/server multilink trunk configuration

Chapter 5: Virtual Link Aggregation Control Protocol

This section provides the concepts and procedures you need to configure the Virtual Link Aggregation Control Protocol (VLACP).

Virtual Link Aggregation Control Protocol

Use Virtual Link Aggregation Control Protocol (VLACP) as an extension to LACP 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 end-to-end communication. After 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.

You can also use VLACP with MLT to complement its capabilities and provide quick failure detection. Avaya recommends VLACP for all SMLT access links after you configure the links as MLT to ensure both end devices are able to communicate.

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 9: Problem description \(1 of 2\)](#) on page 68, after 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.

The following illustration shows an MLT running with VLACP. VLACP can operate end-to-end, but you can also use it as a point-to-point link.




Figure 9: Problem description (1 of 2)

In the following illustration, 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 10: Problem description (2 of 2)

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

If used in conjunction with SMLT, VLACP switches 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

Virtual Services Platform 9000 software uses the following VLACP timers:

- fast periodic timer—100 to 20 000 ms; default 200 ms
- slow periodic timer—10 000 to 30 000 ms; default 30 000 ms

Virtual Services Platform 9000 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.

Though you can configure Virtual Services Platform 9000 ports 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 to achieve sub-100-ms failover. In addition, Avaya only supports the same minimum fast periodic timer configuration of 3 times 30 ms after you enable High Availability mode, no matter how scaled the environment.

VLACP considerations

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

- Avaya recommends that you use VLACP with long timers on MLT IST links.
- If a VLACP-enabled port does not receive a VLACP Data Unit (VLACPDU), it must enter the disabled state. There are occasions where 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. You must either enable or disable both sides of the point-to-point connection.

You can configure VLACP on each port. The port can be either an individual port or an MLT member. VLACPDUs can be sent periodically on each port where VLACP is enabled to exchange VLACPDUs from an end-to-end perspective. If VLACPDUs are not received on a particular link, that link is 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 stay in a disabled state. After you enable VLACP at the IST peer, the VLACPDU is received and the ports are reenabled. The system can replicate this behavior despite the IST connectivity between the end-to-end peers. After 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 stays active until the expiry timeout occurs on the other end. As soon you enable VLACP at the other end, the VLACPDU is received by the peer and the ports are brought up at the software level.

VLACP configuration using ACLI

Configure Virtual LACP (VLACP) to implement link status control protocol at the port level. VLACP detects end-to-end failures in the Avaya Virtual Services Platform 9000. Virtual LACP cannot interoperate with Link Aggregation Control Protocol (LACP).

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. After Hello packets are not received, VLACP transitions to a failure state and disables the port.

About this task

Important:

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

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
```

```
interface GigabitEthernet {slot/port[-slot/port] [, ...]}
```

2. Configure optional parameters for the port. If you do not configure these parameters, the system uses the default values.

- a. Configure the protocol identification for the port:

```
vlacp ethertype <1536-65535 | 0x600-0xffff> [funcmac-addr
0x00:0x00:0x00:0x00:0x00:0x00]
```

- b. Configure the fast or slow periodic times:

```
vlacp fast-periodic-time <100-20000> | slow-periodic-time
<10000-30000>
```

You can configure both parameters in the same command entry.

- c. Configure the timeout parameters:

```
vlacp timeout <long|short> timeout-scale <2-10>
```

You can configure both parameters in the same command entry.

- d. Exit GigabitEthernet Interface Configuration mode:

```
exit
```

- e. Enable VLACP sequence numbers:

```
vlacp sequence-num
```

3. Enable VLACP on a port:

```
vlacp enable
```

Example

Configure VLACP on port 4/2:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface GigabitEthernet 4/2
Switch:1(config-if)#vlacp fast-periodic-time 400 timeout short
Switch:1(config-if)#vlacp enable
```

Variable definitions

Use the data in the following table to help you use the `vlacp` command.

Variable	Value
enable	Enables VLACP for this port. The default is disabled.
etherstype <1536-65535 0x600-0xffff>	Configures the VLACP protocol identification for this port. Enter the type in decimal or hexadecimal format. The default is 0x8103.
fast-periodic-time <100-20000>	Configures the fast periodic time (in milliseconds) for this port. The default is 200.

Table continues...

Variable	Value
funcmac-addr <0x00:0x00:0x00:0x00:0x00:0x00>	Configures the multicast MAC address used for the VLACPDU. Specify a MAC address in the format 0x00:0x00:0x00:0x00:0x00:0x00.
sequence-num	Enables VLACP sequence-num. The default is enabled. If you want to disable sequence-num, ensure you are in Global Configuration mode.
slow-periodic-time <10000-30000>	Configures the slow periodic time (in milliseconds) for a specific port type. The default is 30,000.
timeout {long short}	Configures the port to use the long or short timeout: <ul style="list-style-type: none"> 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. <p>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.</p> <p>The default is long.</p>
timeout-scale <2-10>	Configures a timeout scale for this port used to calculate the timeout. The default value is 3.

Viewing the VLACP port configuration

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

Procedure

View the VLACP port configuration for all interfaces:

```
show vlacp interface gigabitethernet [{slot/port[-slot/port][, ...]] [vid <1-4084>]
```

Example

```
VSP-9012:1#show vlacp interface gigabitethernet 4/1-4/5
```

```
=====
                          VLACP Information
=====
INDEX  ADMIN   OPER    PORT   FAST   SLOW   TIMEOUT  TIMEOUT  ETHER   MAC
      ENABLED ENABLED STATE  TIME   TIME   TIME     SCALE   TYPE   ADDR
-----
4/1   false   false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/2   true    true    UP     400    30000  short    3       0x8103  01:80:c2:
00:11:00
4/3   false   false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/4   false   false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
=====
```

```
00:11:00
4/5 false false DOWN 200 30000 long 3 0x8103 01:80:c2:
00:11:00
```

Variable definitions

Use the data in the following table to use the `show vlacp interface gigabitethernet` command.

Variable	Value
{slot/port[slot/port][,...]}	Specifies a port or list of ports to show only the VLACP information for those ports.
vid <1-4084>	Specifies a VLAN ID to show only ports attached to a particular VLAN. The range is 1–4084.

Enabling or disabling VLACP globally

Use VLACP as an extension to LACP for end-to-end failure detection. Enable or disable VLACP globally to reset the port level configuration. The default is disabled.

About this task

Important:

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

Procedure

1. Log on to the Global Configuration mode:

```
enable
configure terminal
```

2. Enable VLACP globally:

```
vlacp enable
```

3. Disable VLACP globally:

```
no vlacp enable
```

Example

Enable VLACP globally:

```
VSP-9012:1(config)#vlacp enable
```

VLACP configuration using EDM

Configure Virtual LACP (VLACP) to implement link status control protocol at the port level. VLACP cannot interoperate with Link Aggregation Control Protocol (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. After Hello packets are not received, the VLACP transitions to a failure state and the port is disabled.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **VLACP Global** tab.
4. Select the **VlACPEnable** check box.
5. Click **Apply**.

VLACP Global field descriptions

Use the data in the following table to use the **VLACP Global** tab.

Name	Description
VlACPEnable	Enables VLACP globally. The default is disabled.

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

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the **VLACP** tab.
5. Select the **AdminEnable** check box.
6. Configure the remaining parameters as required.
7. Click **Apply**.

VLACP field descriptions

Use the data in the following table to use the **VLACP** tab.

Name	Description
AdminEnable	Enables VLACP for the port. The default is disabled.
OperEnable	Displays 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. Set 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. Set all LACP enabled ports the same value from this setting. The range is 10000–30000. The default value is 30000.
Timeout	Specifies 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	Specifies the VLACP protocol identification. The ID is in hexadecimal format. The default value is 0x8103.
EtherMacAddress	Specifies the multicast MAC address exclusively used for VLACPDUs.
PortState	Displays the VLACP port state.

Glossary

actor	An LACP device at the near end of a link, which performs the action. The device at the far end of the link is called the partner. LACP compares the actor and partner information to determine what action to perform.
Bridge Protocol Data Unit (BPDU)	A data frame used to exchange information among the bridges in local or wide area networks for network topology maintenance.
denial-of-service (DoS)	Attacks that prevent a target server or victim device from performing its normal functions through flooding, irregular protocol sizes (for example, ping requests aimed at the victim server), and application buffer overflows.
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.
Interior Gateway Protocol (IGP)	Distributes routing information between routers that belong to a single Autonomous System (AS).
Internet Control Message Protocol (ICMP)	A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.
interswitch trunking (IST)	A feature that uses one or more parallel point-to-point links to connect two aggregation switches. The two aggregation switches use this channel to share information and operate as a single logical switch. Only one interswitch trunk can exist on each Split Multilink Trunking (SMLT) aggregation switch.
Layer 2	Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
Layer 3	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).
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 increases the redundancy for higher availability.
load balancing	The practice of splitting communication into two (or more) routes or servers.
media	A substance that transmits data between ports; usually fiber optic cables or category 5 unshielded twisted pair (UTP) copper wires.
Network Interface Card (NIC)	A network interface device (NID) in the form of a circuit card installed in an expansion slot of a computer to provide network access.
packet loss	Expressed as a percentage of packets dropped over a specified interval. Keep packet loss to a minimum to deliver effective IP telephony and IP video services.
Protocol Data Units (PDUs)	A unit of data that is specified in a protocol of a specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.
Simple Loop Prevention Protocol (SLPP)	Simple Hello Protocol that prevents loops in a Layer 2 network (VLAN).
Small Form Factor Pluggable (SFP)	A hot-swappable input and output enhancement component used with Avaya products to allow gigabit Ethernet ports to link with other gigabit Ethernet ports over various media types.
SMLT aggregation switch	One of two IST peer switches that form a split link aggregation group. It connects to multiple wiring closet switches, edge switches, or customer premise equipment (CPE) devices.
SMLT client	A switch located at the edge of the network, such as in a wiring closet or CPE connecting to two SMLT aggregation switches. An SMLT client switch performs link aggregation but does not require Split MultiLink Trunking (SMLT) intelligence.
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.
Spanning Tree Group (STG)	A collection of ports in one spanning-tree instance.

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.
trunk	A logical group of ports that behaves like a single large port.
trunk port	A port that connects to the service provider network such as the MPLS environment.
User Datagram Protocol (UDP)	In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/IP host systems use UDP for application-to-application programs.
Virtual Link Aggregation Control Protocol (VLACP)	Virtual Link Aggregation Control Protocol (VLACP) is a Layer 2 handshaking protocol that can detect end-to-end failure between two physical Ethernet interfaces.
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.
wiring closet	A central termination area for telephone or network cabling or both.