



Configuring Link Aggregation, MLT, SMLT and vIST for VOSS

Release 8.2 (VOSS)
9036551-00 Rev AA
July 2020

© 2017-2020, Extreme Networks, Inc.
All Rights Reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see:
www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>

Contents

Chapter 1: About this Document	6
Purpose.....	6
Conventions.....	6
Text Conventions.....	7
Documentation and Training.....	9
Getting Help.....	9
Providing Feedback.....	10
Chapter 2: New in this Document	11
Notice about Feature Support.....	11
Chapter 3: Link Aggregation Control Protocol	12
Link aggregation overview.....	12
Virtual Inter-Switch Trunk (vIST).....	13
MultiLink Trunking with LACP.....	15
Input/output port redundancy.....	18
LACP configuration considerations.....	18
Important Information and Restrictions.....	19
LACP with Simplified vIST/SPB NNI links.....	19
vIST VLAN IP addresses.....	20
Simplified vIST and egress port-based filters.....	20
LACP configuration using CLI.....	20
Configuring global LACP parameters.....	21
Configuring LACP on a port.....	22
Configuring LACP on an MLT.....	25
Configuring LACP and Private VLANs.....	26
Configuring LACP on a VLAN.....	27
Viewing LACP configuration information.....	28
LACP configuration using EDM.....	32
Configuring global LACP parameters.....	32
Configuring LACP parameters.....	34
Configuring LACP on a port.....	36
Configuring LACP on an Extreme Integrated Application Hosting Port.....	41
Chapter 4: MultiLink Trunking and Split MultiLink Trunking	48
Link aggregation overview.....	48
Virtual Inter-Switch Trunk (vIST).....	49
Simplified Virtual-IST.....	51
MultiLink Trunking.....	51
Split MultiLink Trunking.....	54
SLPP Guard.....	60
MLT and SMLT configuration considerations.....	61

MLT and Private VLANs.....	64
MLT and SMLT link aggregation configuration using the CLI.....	65
Configuring MLT.....	65
Viewing MLT port members.....	67
Adding ports to an MLT LAG.....	68
Removing ports from an MLT LAG.....	69
Creating an SMLT from an existing MLT.....	70
Virtual interswitch trunk (vIST).....	71
Viewing all ports configured for SMLT.....	77
Viewing information about collision errors.....	78
Viewing information about Ethernet errors.....	79
SLPP Guard configuration.....	79
MLT and SMLT Link Aggregation Configuration using EDM.....	82
Adding a multilink or LACP trunk.....	83
Adding ports to an MLT.....	85
Viewing trunks.....	86
Creating a virtual IST using EDM.....	86
Editing a virtual IST.....	87
Configuring Simplified vIST in SMLT topologies.....	88
SLPP Guard configuration.....	90
MLT configuration examples.....	93
MultiLink Trunking.....	93
MultiLink Trunking with Link Aggregation Control Protocol.....	94
MLT network topology and configuration reference.....	95
Example 1: Switch-to-switch MLT.....	95
Example 2: Switch-to-server MLT.....	96
Example 3: Client/Server MLT.....	96
Chapter 5: Virtual Link Aggregation Control Protocol.....	98
Virtual Link Aggregation Control Protocol.....	98
VLACP Flap Detect and Damping.....	100
VLACP Configuration using CLI.....	101
Configure VLACP on a port.....	102
View the VLACP Port Configuration.....	103
Enabling or disabling VLACP globally.....	105
Configure VLACP Flap Detect and Damping on a Port.....	106
VLACP Configuration using EDM.....	107
Enabling VLACP globally.....	108
Configure VLACP on a port.....	108
Configuring VLACP on an Extreme Integrated Application Hosting Port.....	110
Chapter 6: Link-state tracking (LST).....	112
Link-state tracking (LST) Overview.....	112
LST configuration using CLI.....	114
Configuring LST.....	114

LST configuration using EDM.....	116
Configuring LST.....	116
Glossary	118

Chapter 1: About this Document

This section discusses the purpose of this document, the conventions used, ways to provide feedback, additional help, and information regarding other Extreme Networks publications.

Purpose

This document provides information on features in VSP Operating System Software (VOSS). VOSS runs on the following product families:

- Extreme Networks VSP 4000 Series (includes VSP 4450 Series)
- Extreme Networks VSP 4900 Series
- Extreme Networks VSP 7200 Series
- Extreme Networks VSP 7400 Series
- Extreme Networks VSP 8000 Series (includes VSP 8200 Series and VSP 8400 Series)
- Extreme Networks VSP 8600 Series
- Extreme Networks XA1400 Series

 **Note:**

VOSS is licensed on the XA1400 Series as a Fabric Connect VPN (FCVPN) application, which includes a subset of VOSS features. FCVPN transparently extends Fabric Connect services over third-party provider networks.

This document contains conceptual, procedural information and instructions to help you use, configure and manage link aggregation and MultiLink Trunking on the VOSS switches.

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Conventions

This section discusses the conventions used in this guide.

Text Conventions

The following tables list text conventions that can be used throughout this document.

Table 1: Notice Icons







Icon	Alerts you to...
 Important:	A situation that can cause serious inconvenience.
 Note:	Important features or instructions.
 Tip:	Helpful tips and notices for using the product.
 Danger:	Situations that will result in severe bodily injury; up to and including death.
 Warning:	Risk of severe personal injury or critical loss of data.
 Caution:	Risk of personal injury, system damage, or loss of data.

Table 2: Text Conventions

Convention	Description
Angle brackets (< >)	<p>Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command.</p> <p>If the command syntax is <code>cfm maintenance-domain maintenance-level <0-7></code>, you can enter <code>cfm maintenance-domain maintenance-level 4</code>.</p>
Bold text	<p>Bold text indicates the GUI object name you must act upon.</p> <p>Examples:</p> <ul style="list-style-type: none"> • Click OK. • On the Tools menu, choose Options.
Braces ({ })	<p>Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command.</p> <p>For example, if the command syntax is <code>ip address {A.B.C.D}</code>, you must enter the IP address in dotted, decimal notation.</p>

Table continues...

Convention	Description
Brackets ([])	<p>Brackets ([]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command.</p> <p>For example, if the command syntax is <code>show clock [detail]</code>, you can enter either <code>show clock</code> or <code>show clock detail</code>.</p>
Ellipses (...)	<p>An ellipsis (...) indicates that you repeat the last element of the command as needed.</p> <p>For example, if the command syntax is <code>ethernet/2/1 [<parameter> <value>]...</code>, you enter <code>ethernet/2/1</code> and as many parameter-value pairs as you need.</p>
<i>Italic Text</i>	<p>Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.</p>
Plain Courier Text	<p>Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages.</p> <p>Examples:</p> <ul style="list-style-type: none"> • <code>show ip route</code> • <code>Error: Invalid command syntax</code> <code>[Failed][2013-03-22 13:37:03.303</code> <code>-04:00]</code>
Separator (>)	<p>A greater than sign (>) shows separation in menu paths.</p> <p>For example, in the Navigation tree, expand the Configuration > Edit folders.</p>
Vertical Line ()	<p>A vertical line () separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command.</p> <p>For example, if the command syntax is <code>access-policy by-mac action { allow deny }</code>, you enter either <code>access-policy by-mac action allow</code> or <code>access-policy by-mac action deny</code>, but not both.</p>

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and software compatibility](#) for Extreme Networks products

[Extreme Optics Compatibility](#)

[Other resources](#) such as white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

[Extreme Portal](#)

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

[The Hub](#)

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

[Call GTAC](#)

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Service Notifications

You can subscribe to email notifications for product and software release announcements, Vulnerability Notices, and Service Notifications.

1. Go to www.extremenetworks.com/support/service-notification-form.
2. Complete the form (all fields are required).
3. Select the products for which you would like to receive notifications.

*** Note:**

You can modify your product selections or unsubscribe at any time.

4. Select **Submit**.

Providing Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information in the document.
- Broken links or usability issues.

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Chapter 2: New in this Document

There are no feature changes in this document.

Notice about Feature Support

This document includes content for multiple hardware platforms across different software releases. As a result, the content can include features not supported by your hardware in the current software release.

If a documented command, parameter, tab, or field does not display on your hardware, it is not supported.

For information about physical hardware restrictions, see your hardware documentation.

Chapter 3: Link Aggregation Control Protocol

The following sections provide 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.

Related links

- [MultiLink Trunking with LACP](#) on page 15
- [LACP configuration considerations](#) on page 18
- [LACP configuration using CLI](#) on page 20
- [LACP configuration using EDM](#) on page 32

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 switch 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.
- 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 switch software offers LACP functionality layered with MLT. This document uses the term MLT with LACP to refer to this functionality.

Split MultiLink Trunking (SMLT)

Split MultiLink Trunking (SMLT) is an option that improves Layer 2 (bridged) resiliency by providing for the addition of 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 virtualized Switch Cluster through the SPBM cloud and are referred to as a Virtual Inter-Switch Trunk (vIST) Core Switch pair.

For more information on SMLT, see [Split MultiLink Trunking](#) on page 54.

LACP with SMLT

You can use LACP on SMLT configurations. The switch provides modifications to the LACP in SMLT configurations. LACP-capable devices can connect to an SMLT aggregation pair.

*** Note:**

Virtual IST is not supported on LACP-enabled MLTs.

VLACP with SMLT

You can also configure Virtual LACP (VLACP) with an SMLT configuration. VLACP is a modification that provides end-to-end failure detection. VLACP is not a link aggregation protocol.

VLACP implements link status control protocol at the port level. 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.

*** Note:**

Do not configure VLACP on LACP-enabled ports. VLACP does not operate properly with LACP.

Virtual Inter-Switch Trunk (vIST)

Table 3: Virtual Inter-Switch Trunk product support

Feature	Product	Release introduced
For configuration details, see Configuring Link Aggregation, MLT, SMLT and vIST for VOSS .		
Switch cluster (multi-chassis LAG) -Virtual Inter-Switch Trunk (vIST)	VSP 4450 Series	VOSS 4.1
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 6.1
	XA1400 Series	Not Supported

Split MultiLink Trunking provides subsecond failover when a switch fails. Virtual Inter-Switch Trunk (vIST) improves upon that Layer 2 and Layer 3 resiliency by using a virtualized IST channel through the SPBM cloud. The vIST channel carries the vIST control traffic and data traffic during an SMLT failover. This feature dramatically improves resiliency over other methods.

Because vIST uses a virtual channel and because IS-IS runs over it, vIST eliminates the potential single point of failure with a dedicated MLT. The vIST channel is always up as long as there is SPBM connectivity between the vIST peers.

vIST interoperates between any two devices that support vIST, and the devices do not have to be of the same type.

vIST creates a virtualized channel through the SPBM cloud, and this channel connects two SMLT devices to form a virtualized Switch Cluster. Note that the SPBM cloud can consist of as few as two nodes.

 **Important:**

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of the header size increase when transmitting packets over the SPBM cloud.

 **Note:**

- Users may observe a momentary increase in activity when a MAC delete message is received from a peer. This is due to vIST engaging in MAC learning activities. This is a normal operational procedure.
- For proper traffic flow, if a Layer 2 VSN is created on one vIST peer, it must also be created on the other vIST peer. For more information on Layer 2 VSN, see [Configuring Fabric Basics and Layer 2 Services for VOSS](#).

vIST configuration note

If you need to update the vIST VLAN IP address on vIST peers by deleting and recreating the vIST vlan IP address (for example, as part of maintenance), ensure that you update one vIST BEB at a time.

 **Caution:**

Always perform vIST configuration updates under no traffic. Otherwise, it results in traffic loss.

Before you begin updating a device, as a first step, isolate the device by shutting down all the links and failing over the traffic to its vIST peer. Then, delete and recreate the vIST VLAN IP on the device and save your configuration. When bringing the device back into operation, first unshut those NNI ports that bring up vIST, followed by the SMLT configured ports, and then all the remaining ports, to prevent network loops or duplicate traffic.

For information on vIST configuration, see [Creating a Virtual IST](#) on page 71 or [Creating a virtual IST using EDM](#) on page 86.

vIST operational note

When you enable IST and boot the chassis, the SMLT enabled trunk ports (SMLT ports) are automatically locked. A timeout mechanism automatically unlocks the SMLT ports when the IST control channel fails to establish within a reasonable amount of time. The timeout mechanism prevents the SMLT ports from being locked forever. Initially 240 seconds are allowed for the switch to determine the IST VLAN status.

The IST VLAN is considered up if at least one port is forwarding traffic and an ARP entry is populated for the IP address of the IST peers. Once the IST VLAN is up, the timeout value is reset to 60 seconds. The IST control channel must be up within the timeout period. If the timeout period is

exceeded, then the SMLT ports are automatically unlocked and a message is logged stating that the SMLT ports are unlocked due to a timeout.

*** Note:**

If the IST filter is enabled before the timeout, then the IST filter is unaffected and remains enabled.

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.

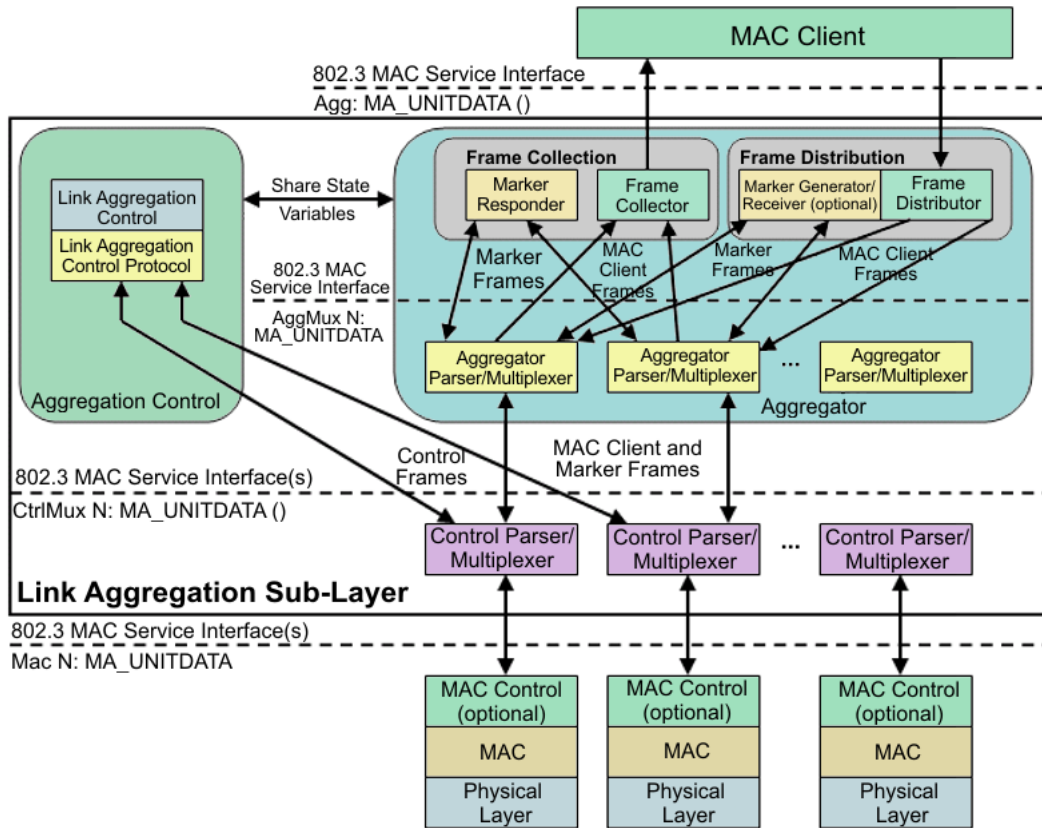


Figure 1: Link aggregation sublayer (according to IEEE 802.3ad)

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. The switch only implements the Marker Receiver function. For more information about the frame distribution algorithm, see [MLT traffic distribution algorithm](#) on page 52.
- 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.
- 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.

Related links

[Link Aggregation Control Protocol](#) on page 12

[Input/output port redundancy](#) on page 18

Input/output port redundancy

You can use the MLT link aggregation mechanism to protect I/O ports. MLT is compatible with 802.3ad static, and provides a load sharing and failover mechanism to protect against module, port, fiber, or complete link failures.

You can use MLT with Link Access Control Protocol (LACP) disabled or use LACP enabled by itself.

Related links

[MultiLink Trunking with LACP](#) on page 15

LACP configuration considerations

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

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.
- System priority—generates the switch ID after communicating with other systems. It is recommended 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 must 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.

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.

The switch 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.

LACP and private VLANs

- When using LACP, configure the private-vlan at the interface level.

Link aggregation scaling

For the latest applicable scaling information, see [Release Notes for VOSS](#) for the version of the software running on the switch.

Important Information and Restrictions

Before you use the switch, review the following important information and restrictions.

LACP with Simplified vIST/SPB NNI links

LACP is not recommended on SPB NNI MLT links or on the Simplified Virtual IST.

vIST VLAN IP addresses

Do not configure a Rendezvous Point (RP) or Bootstrap Router (BSR) on the vIST VLAN because you cannot ping them outside of the vIST VLAN subnet. When you enter the `ip pim enable` command on the vIST VLAN, the following message displays:

```
WARNING: Please do not use virtual IST VLAN IP address for BSR and RP
related configurations, as unicast packets to virtual IST vlan IP address
from outside of virtual IST vlan subnet will be dropped. Use Loopback or
CLIP interface IP address for BSR and RP related configurations.
```

Simplified vIST and egress port-based filters

On Simplified vIST nodes, egress port-based filters may not work for IP multicast routed traffic because vIST internal filter rules (to prevent duplicate traffic) have higher precedence than user-created filters.

 **Note:**

- This issue is specific to IP multicast routed traffic only.
- Egress port-based filters work for Layer 2 multicast, broadcast and unicast traffic.

LACP configuration using CLI

This section describes how to configure and manage link aggregation using the Command Line Interface (CLI), 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.

- The switch does not support standby ports for LACP aggregation groups.

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.

A change to the global parameter configuration takes effect after you restart the LACP globally or on each port.

About this task

! Important:

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

Procedure

1. Enter 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> ] [system-priority <0-65535> ]
[timeout-scale <2-10>]
```

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

Example

```
Switch:1(config)#lacp fast-periodic-time 2000
Switch: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 value is 2000.
enable	Enables LACP globally. The default value is disabled.
fast-periodic-time <200–20000>	Configures the fast periodic time (in milliseconds) globally. The default value is 1000.
slow-periodic-time <10000–30000>	Configures the slow periodic time globally. The default value 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 LACP system priority globally. The default value 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.

You must 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.

Note:

When enabling or disabling LACP on a port, it is recommended to disable the port first and re-enable the port after the configuration is complete.

The minimum LACP configuration is as follows:

- Assign a given key to a set of ports.
- Assign the same key to an MLT with no members. The ports will automatically become MLT members.

Keys do not have to match between two LACP peers.

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.

Note:

When using LACP with private VLANs, configure the private VLAN at the interface level.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-
port]][, ...]} or interface vlan <1-4059>
```

*** Note:**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Assign an LACP key to a set of ports.

```
lacp key <1-512|defVal>
```

3. Change the LACP mode:

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

4. Change the port priority:

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

5. **(Optional)** Change the system priority:

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

6. 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.

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

*** Note:**

All the parameters beginning with **partner** are for debug purposes only. If you configure these commands locally and there is a mismatch with what's learned from the partner, there will be trace or log message.

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

8. Configure additional LACP parameters as required:

```
lacp enable [fast-periodic-time <200-20000>] [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 1/2 and 1/3:

```
Switch:1(config)# interface gigabitethernet 1/2-1/3
Switch:1(config-if)# lacp key 1 timeout short
Switch:1(config-if)# lacp aggregation enable
Switch:1(config-if)# lacp enable
```

Variable definitions

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

* Note:

All the parameters beginning with **partner** are for debug purposes only. If you configure these commands locally and there is a mismatch with what's learned from the partner, there will be trace or log message.

Variable	Value
aggr-wait-time <200–2000>	Configures the aggregation wait time (in milliseconds) for this port. The default is 2000.
aggregation enable	Enables aggregation on the port, which makes it an aggregated link.
enable	Enables LACP for this port. The default is disabled.
fast-periodic-time <200–20000>	Configures the fast periodic time (in milliseconds) for this port. The default is 1000 ms.
key <1-512 defVal>	Configures the aggregation key for this port. Enter the aggregation key value or defVal (1024 + IfIndex)
mode {active passive}	Configures the LACP mode to be active or passive.
partner-key <0–65535>	Configures the partner administrative key.
partner-port <0–65535>	Configures the partner administrative port value.
partner-port-priority <0–65535>	Configures the partner administrative port priority value.
partner-state <0-255 0x0-0xff>	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. For example, to set the two partner-state parameters <ul style="list-style-type: none"> • Act = true • Agg = true

Table continues...

Variable	Value
	specify a value of 0x05 (bitmap = 00000101).
partner-system-id <0x00:0x00:0x00:0x00:0x00:0x00>	Configures the partner administrative system ID. Specify a MAC address in the format 0x00:0x00:0x00:0x00:0x00:0x00.
partner-system-priority <0–65535>	Configures the partner administrative system priority value.
priority <0–65535>	Configures the port priority. The default value is 32768. To set this option to the default value, use the default operator with the command.
slow-periodic-time <10000-30000>	Configures the slow periodic time for this port. The default is 30000 ms. To set this option to the default value, use the default operator with the command.
system-priority <0-65535>	Configures the system priority for this port. The default is 32768.
timeout-scale <2-10>	Configures a timeout scale for this port. The default value is 3. 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.
timeout-time {long short}	Configures the timeout to either long or short.

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. When you add the VLAN to the MLT, also add the inactive aggregation ports to the VLAN. You may need 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. Enter 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:

```
Switch:1(config)# interface mlt 3
Switch: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.

Configuring LACP and Private VLANs**About this task**

Use the following procedure to configure Link Aggregation Control Protocol (LACP) on a private VLAN.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]][, ...]} or interface vlan <1-4059>
```

*** Note:**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable trunking:

```
encapsulation dot1q
```

3. Set private VLAN port type:

```
private-vlan [isolated|promiscuous|trunk]
```

Example

```
Switch:1(config)#interface GigabitEthernet 1/6
Switch:1(config-if)#encapsulation dot1q
Switch:1(config-if)#private-vlan promiscuous
Switch:1(config-if)#exit
```

```
Switch:1(config)#interface GigabitEthernet 1/37
Switch:1(config-if)#encapsulation dot1q
Switch:1(config-if)#private-vlan promiscuous
```

Configuring LACP on a VLAN

Use this procedure to configure LACP on a VLAN.

About this task**! Important:**

Configuration at the global level overrides and resets all settings at the VLAN level.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
configure terminal
interface vlan <1-4059>
```

2. Configure global aggregation wait time for the VLAN:

```
lacp agrgr-wait-time <200-2000>
```

3. Enable LACP:

```
lacp enable
```

4. Configure the global periodic time:

```
lacp enable [fast-periodic-time <200-20000>] [slow-periodic-time
<10000-30000>]
```

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

5. Configure a global LACP system ID.

```
lacp smlt-sys-id 0x00:0x00:0x00:0x00:0x00:0x00
```

6. Configure the system priority:

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

7. Configure the timeout-scale:

```
lacp timeout-scale <2-10>
```

Example

```
Switch:1(config)#interface vlan 20
```

```
Switch:1(config-if)#lacp fast-periodic-time 2000
```

```
Switch:1(config-if)#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 for the VLAN. The default is disabled.
fast-periodic-time <200–20000>	Configures the fast periodic time (in milliseconds) globally for the VLAN. The default is 1000 ms.
slow-periodic-time <10000-30000>	Configures the slow periodic time globally for the VLAN. The default is 30000 ms.
smlt-sys-id 0x00:0x00:0x00:0x00:0x00:0x00	Configures the LACP system ID globally.
system-priority <0-65535>	Configures the LACP system priority globally.
timeout-scale <2-10>	Configures a timeout scale for the VLAN. The default value is 3. 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.

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 local device:

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

OR

```
show lacp actor-admin interface gigabitethernet [vid <1-4059>]
[{{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]]
```

3. View LACP operational information for the actor device:

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

OR

```
show lacp actor-oper interface gigabitethernet [vid <1-4059>]
[{{slot/port[/sub-port]}[-slot/port[/sub-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 <1-4059>] [{{slot/port[/sub-
port]}[-slot/port[/sub-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 <1-4059>]
[{{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]]
```

7. View LACP operational information for the partner device:

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

OR

```
show lacp partner-oper interface gigabitethernet [vid <1-4059>]
[{{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]]
```

Example

```
Switch: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.

```

Switch:1(config-if)# show lacp actor-admin interface gigabitethernet 1/1
=====
                        Actor Admin
=====
INDEX SYS   SYS           KEY   PORT   PORT   STATE
   Prio  ID                                     Prio
-----
1/1  32768 00:24:7f:9f:c0:00 1     0x114   32768 act short aggr
=====
    
```

```

Switch:1(config-if)# show lacp partner-admin interface gigabitethernet 1/1
=====
                        Partner Admin
=====
INDEX SYS   SYS           KEY   PORT   PORT   STATE
   Prio  ID                                     Prio
-----
1/1  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
<code>actor-admin interface gigabitethernet [vid <1-4059> {slot/port[/sub-port] [-slot/port[/sub-port]] [...]}]</code>	<p>Shows LACP actor (or local) administrative information for all interfaces or the specified interface.</p> <ul style="list-style-type: none"> Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. <code>{slot/port[/sub-port] [-slot/port[/sub-port]] [...]}</code> is the port or port list.

Table continues...

Variable	Value
actor-oper interface gigabitethernet [vid <1-4059> [{slot/port[/sub-port] [-slot/port[/sub-port]] [...]}]	Shows LACP actor operational information for all interfaces or the specified interface. <ul style="list-style-type: none"> • Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • {slot/port[/sub-port] [-slot/port[/sub-port]] [...]} is the port or port list.
extension interface gigabitethernet [vid <1-4059> [{slot/port[/sub-port] [-slot/port[/sub-port]] [...]}]	Shows LACP timer information for all interfaces or the specified interface. <ul style="list-style-type: none"> • Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • {slot/port[/sub-port] [-slot/port[/sub-port]] [...]} is the port or port list.
interface [gigabitethernet [vid <1-4059> [{slot/port[/sub-port] [-slot/port[/sub-port]] [...]}] [mlt <64-6399>] [mlt id <1-128>]]	Shows all LACP port configuration information for all interfaces or the interface you specify. <ul style="list-style-type: none"> • Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • {slot/port[/sub-port] [-slot/port[/sub-port]] [...]} is the port or port list. • <64-6399> is the interface index of the mlt. • <1-128> is the MLT ID.

Table continues...

Variable	Value
partner-admin interface[gigabitethernet] [vid <1-4059> [{slot/port[/sub-port] [-slot/port[/sub-port]] [...]]]	Shows LACP partner administrative information for all interfaces or the specified interface. <ul style="list-style-type: none"> Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. {slot/port[/sub-port] [-slot/port[/sub-port]] [...]} is the port or port list.
partner-oper interface [gigabitethernet] [vid <1-4059> [{slot/port[/sub-port] [-slot/port[/sub-port]] [...]]]	Shows LACP partner operational information for all interfaces or the specified interface. <ul style="list-style-type: none"> Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. {slot/port[/sub-port] [-slot/port[/sub-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.

*** Note:**

The switch does not support standby ports for LACP aggregation groups.

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 tree, 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.

Table continues...

Name	Description
	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.
SysId	Specifies the LACP system ID. The default value is f8:73:a2:00:90:00.
SmltSysId	Specifies the LACP system ID for Split MultiLink Trunking (SMLT).

Configuring LACP parameters

Configure LACP parameters to manage LACP information.

About this task

Important:

The switch does not support standby ports for LACP aggregation groups.

Procedure

1. In the navigation tree, 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 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	Specifies the current administrative value of the key for the aggregator, which is a 16-bit read-write value. The administrative key value can differ from the operational key value. This key needs to match the LAG key. The default value is 0.
ActorOperKey	Displays the current read-only operational value of the key for the aggregator. The operational key value can differ from the administrative key 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.

Table continues...

Name	Description
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.

You must 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 that match the MLT key into that multilink trunk.

The minimum LACP configuration is as follows:

- Assign a given key to a set of ports.
- Assign the same key to an MLT with no members. The ports will automatically become MLT members.

Keys do not have to match between two LACP peers.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation tree, 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. To assign an LACP key, configure the **ActorAdminKey** field.
7. Configure the remaining parameters as required.
8. Click **Apply**.

LACP field descriptions

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

All the parameters beginning with **Partner** are for debug purposes only. If you configure these commands locally and there is a mismatch with what is learned from the partner, there will be trace or log message.

Name	Description
AdminEnable	Enables LACP status for the port. The default value is false.
OperEnable	Displays the operational status of LACP for the port.

Table continues...

Name	Description
	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.
ActorSystemPriority	Specifies the two octet read-write value indicating the priority value associated with the actor system ID. The range is 0–65535. The default value is 32768.
ActorSystemID	Displays the six octet read-write MAC address value used as a unique identifier for the system that contains this aggregator. From the perspective of the link aggregation mechanisms, only a single combination of actor system ID and system priority are considered, and no

Table continues...

Name	Description
	<p>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 read-write administrative value of the key for the aggregator, which is a 16-bit value in the range of 0–65535. The administrative key value can differ from the operational key value. This key needs to match the LAG key.</p>
ActorOperKey	<p>Displays the current read-only operational value of the key for the aggregator. The operational key value can differ from the administrative key 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>
ActorPort	<p>Displays the port number locally assigned to the aggregation port. The port number is communicated in LACPDU as the Actor_Port. This value is read-only.</p>
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 LACPDU, 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)

Table continues...

Name	Description
	<ul style="list-style-type: none"> • 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 LACPDU. 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	<p>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 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</p>

Table continues...

Name	Description
	<p>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 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</p>

Table continues...

Name	Description
	contain the manually configured value carried in PartnerAdminPortPriority if there is no protocol partner. This 16 bit value is read-only. The default value is 0.
PartnerAdminState	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. The default value is none.
PartnerOperState	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. The default value is none.

Configuring LACP on an Extreme Integrated Application Hosting Port

About this task

Perform this procedure to configure Link Aggregation Control Protocol (LACP) on an Extreme Integrated Application Hosting (IAH) port. You must use the LACP keys to determine which IAH ports are eligible for link aggregation. The LACP keys are defined by the IAH ports after you configure the multilink trunk (MLT). You can aggregate the IAH port key that matches the MLT key into that multilink trunk.

The minimum LACP configuration is as follows:

- Assign a given key to a set of IAH ports.
- Assign the same key to an MLT with no members. The IAH ports automatically become MLT members.

The keys do not have to match between two LACP peers.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Insight Port**.
2. Select the IAH port you want to configure.
3. Select the **LACP** tab.
4. Select **AdminEnable**.
5. In the **ActorAdminKey** field, enter a value.

6. Configure the other parameters as required.
7. Click **Apply**.

LACP Field Descriptions

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

*** Note:**

The Partner fields in the LACP tab are for debug purpose only. If you configure them locally and there is a mismatch with what is learned from the partner, trace or log messages are generated.

Name	Description
AdminEnable	Enables LACP status for the Extreme Integrated Application Hosting port. The default value is false.
OperEnable	Specifies the LACP operational status for the Extreme Integrated Application Hosting port. The default value is false.
FastPeriodicTime	Specifies the number of milliseconds between periodic transmissions using short timeouts for all LACP enabled Extreme Integrated Application Hosting ports. The default value is 1000.
FastPeriodicTimeOper	Specifies the operating value of the fast periodic timer on the Extreme Integrated Application Hosting port. The default value is 1000.
SlowPeriodicTime	Specifies the number of milliseconds between periodic transmissions using long timeouts for all LACP enabled Extreme Integrated Application Hosting ports. The default value is 30000.
SlowPeriodicTimeOper	Specifies the operating value of the slow periodic timer on the Extreme Integrated Application Hosting port. The default value is 30000.
AggrWaitTime	Specifies the number of milliseconds to delay aggregation to enable multiple links to aggregate simultaneously. The default value is 2000.
AggrWaitTimeOper	Specifies the operating value of the aggregate wait timer on the Extreme Integrated Application Hosting port.

Table continues...

Name	Description
	The default value is 2000.
TimeoutScale	<p>Specifies the value used to calculate timeout duration from the periodic time. Set the same value for all LACP enabled Extreme Integrated Application Hosting ports.</p> <p>The default value is 3.</p>
TimeoutScaleOper	<p>Specifies the operating value of the timeout scale on the Extreme Integrated Application Hosting port.</p> <p>The default value is 3.</p>
ActorSystemPriority	<p>Specifies the two octet read-write value indicating the priority value associated with the actor system ID.</p> <p>The default value is 32768.</p>
ActorSystemID	<p>Specifies 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 Extreme Integrated Application Hosting 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 Extreme Integrated Application Hosting 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 read-write administrative value of the key for the aggregator, which is a 16-bit value. The administrative key value can differ from the operational key value. This key must match the LAG key.</p>
ActorOperKey	<p>Specifies the current read-only operational value of the key for the aggregator. The operational key value can differ from the administrative key value. The meaning of particular key values is of local significance.</p>

Table continues...

Name	Description
SelectedAggID	Specifies the identifier value of the aggregator that this aggregation port has currently selected. Zero indicates that the aggregation port has not selected an aggregator, either because it is in the process of detaching from an aggregator or because there is no suitable aggregator available for it to select. This value is read-only.
AttachedAggID	Specifies the identifier value of the aggregator to which this aggregation port is currently attached. Zero indicates that the aggregation port is not currently attached to an aggregator. This value is read-only.
ActorPort	Specifies 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	Specifies the priority value assigned to this aggregation port. This 16-bit value is read-write. The default value is 32768.
ActorAdminState	Specifies a string of eight bits, corresponding to the administrative values as transmitted by the actor in LACPDUs. The values are: <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) These values enable administrative control over the values of LACP_Activity, LACP_Timeout, and aggregation. This attribute value is read-write.
ActorOperState	Specifies 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. The default value is lacpActive.

Table continues...

Name	Description
PartnerAdminSystemPriority	<p>Specifies the current administrative value of the Extreme Integrated Application Hosting port number for the protocol partner. It is a 16-bit read-write value. 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	<p>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 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. 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>Specifies 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</p>

Table continues...

Name	Description
	<p>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. 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>Specifies the operational Extreme Integrated Application Hosting port number assigned to this aggregation port by the aggregation port protocol partner. The value of this attribute can contain the manually configured value carried in AggPortPartnerAdminPort if there is no protocol partner. This 16-bit value is read-only.</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. 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>Specifies 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>Specifies 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,</p>

Table continues...

Name	Description
	this value can reflect the manually configured value PartnerAdminState. This attribute value is read-only. The default value is none.

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 switch 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.
- 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 switch software offers LACP functionality layered with MLT. This document uses the term MLT with LACP to refer to this functionality.

Split MultiLink Trunking (SMLT)

Split MultiLink Trunking (SMLT) is an option that improves Layer 2 (bridged) resiliency by providing for the addition of 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 virtualized Switch Cluster through the SPBM cloud and are referred to as a Virtual Inter-Switch Trunk (vIST) Core Switch pair.

For more information on SMLT, see [Split MultiLink Trunking](#) on page 54.

LACP with SMLT

You can use LACP on SMLT configurations. The switch provides modifications to the LACP in SMLT configurations. LACP-capable devices can connect to an SMLT aggregation pair.

*** Note:**

Virtual IST is not supported on LACP-enabled MLTs.

VLACP with SMLT

You can also configure Virtual LACP (VLACP) with an SMLT configuration. VLACP is a modification that provides end-to-end failure detection. VLACP is not a link aggregation protocol.

VLACP implements link status control protocol at the port level. 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.

*** Note:**

Do not configure VLACP on LACP-enabled ports. VLACP does not operate properly with LACP.

Virtual Inter-Switch Trunk (vIST)

Table 4: Virtual Inter-Switch Trunk product support

Feature	Product	Release introduced
For configuration details, see Configuring Link Aggregation, MLT, SMLT and vIST for VOSS .		
Switch cluster (multi-chassis LAG) -Virtual Inter-Switch Trunk (vIST)	VSP 4450 Series	VOSS 4.1
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 6.1
	XA1400 Series	Not Supported

Split MultiLink Trunking provides subsecond failover when a switch fails. Virtual Inter-Switch Trunk (vIST) improves upon that Layer 2 and Layer 3 resiliency by using a virtualized IST channel through the SPBM cloud. The vIST channel carries the vIST control traffic and data traffic during an SMLT failover. This feature dramatically improves resiliency over other methods.

Because vIST uses a virtual channel and because IS-IS runs over it, vIST eliminates the potential single point of failure with a dedicated MLT. The vIST channel is always up as long as there is SPBM connectivity between the vIST peers.

vIST interoperates between any two devices that support vIST, and the devices do not have to be of the same type.

vIST creates a virtualized channel through the SPBM cloud, and this channel connects two SMLT devices to form a virtualized Switch Cluster. Note that the SPBM cloud can consist of as few as two nodes.

 **Important:**

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of the header size increase when transmitting packets over the SPBM cloud.

 **Note:**

- Users may observe a momentary increase in activity when a MAC delete message is received from a peer. This is due to vIST engaging in MAC learning activities. This is a normal operational procedure.
- For proper traffic flow, if a Layer 2 VSN is created on one vIST peer, it must also be created on the other vIST peer. For more information on Layer 2 VSN, see [Configuring Fabric Basics and Layer 2 Services for VOSS](#).

vIST configuration note

If you need to update the vIST VLAN IP address on vIST peers by deleting and recreating the vIST vlan IP address (for example, as part of maintenance), ensure that you update one vIST BEB at a time.

 **Caution:**

Always perform vIST configuration updates under no traffic. Otherwise, it results in traffic loss.

Before you begin updating a device, as a first step, isolate the device by shutting down all the links and failing over the traffic to its vIST peer. Then, delete and recreate the vIST VLAN IP on the device and save your configuration. When bringing the device back into operation, first unshut those NNI ports that bring up vIST, followed by the SMLT configured ports, and then all the remaining ports, to prevent network loops or duplicate traffic.

For information on vIST configuration, see [Creating a Virtual IST](#) on page 71 or [Creating a virtual IST using EDM](#) on page 86.

vIST operational note

When you enable IST and boot the chassis, the SMLT enabled trunk ports (SMLT ports) are automatically locked. A timeout mechanism automatically unlocks the SMLT ports when the IST control channel fails to establish within a reasonable amount of time. The timeout mechanism prevents the SMLT ports from being locked forever. Initially 240 seconds are allowed for the switch to determine the IST VLAN status.

The IST VLAN is considered up if at least one port is forwarding traffic and an ARP entry is populated for the IP address of the IST peers. Once the IST VLAN is up, the timeout value is reset to 60 seconds. The IST control channel must be up within the timeout period. If the timeout period is exceeded, then the SMLT ports are automatically unlocked and a message is logged stating that the SMLT ports are unlocked due to a timeout.

 **Note:**

If the IST filter is enabled before the timeout, then the IST filter is unaffected and remains enabled.

Simplified Virtual-IST

Simplified Virtual-IST (vIST) is for conventional switch clustering deployments that use SMLT and not SPB. The Simplified vIST feature provides a single CLI command to enable the virtual IST for SMLT deployments.

- Simplified vIST is available for conventional multicast deployments with PIM and IGMP only when the boot flag (`spbm-config-mode`) is disabled.

! Important:

PIM is supported with Simplified vIST only, not with SPB vIST.

- When the `spbm-config-mode` boot flag is enabled (default setting), Simplified vIST is not available. This means that you continue to configure SPB/IS-IS for vIST as described in [Creating a Virtual IST](#) on page 71 and [Creating a virtual IST using EDM](#) on page 86.
- Simplified VIST requires that the two vIST devices be directly connected.
- For legacy IGMP snooping and IGMP snooping over Simplified vIST, the IGMP sender information is NOT synchronized across the vIST. This means that `show ip igmp sender` displays the sender record only on the switch that the multicast stream actually ingresses.
- In Simplified vIST mode, LACP is not recommended on vIST MLT.

* Note:

For Simplified vIST deployment, if a VLAN is part of an SMLT it must be configured on both the IST peers.

! Important:

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of a header size increase.

For configuration information, see [Configuring Simplified vIST in SMLT topologies](#) on page 75 or [Configuring Simplified vIST in SMLT topologies](#) on page 88.

For information about how to configure Simplified vIST with multicast, see [Configuring IP Multicast Routing Protocols for VOSS](#).

MultiLink Trunking

Table 5: MultiLink Trunking product support

Feature	Product	Release introduced
For configuration details, see Configuring Link Aggregation, MLT, SMLT and vIST for VOSS .		

Table continues...

Feature	Product	Release introduced
MultiLink Trunking (MLT) / Link Aggregation Group (LAG)	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50

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 MLT algorithm ensures that each packet in a flow does not arrive out of sequence, and that a flow always traverses the same link path.

The hashing algorithm uses the following packet fields and the incoming interface (source) port number to calculate the index to outgoing (destination) port number in a MLT:

Traffic type	Hashing algorithm
IPv4 traffic	Hash Key = [Destination IP Address (32 bits), Source IP Address (32 bits), Source TCP/UDP Port, Destination TCP/UDP port]
IPv4 traffic without TCP/UDP header	Hash Key = [Source IP Address (32 bits), Destination IP address (32 bits)]
IPv6 traffic	Hash Key = [Destination IPv6 Address (128 bits), Source IPv6 address (128 bits), Source TCP/UDP Port, Destination TCP/UDP port]
IPv6 traffic without TCP/UDP header	Hash Key = [Source IP Address (128 bits), Destination IP address (128 bits)]
Mac-In-Mac transit traffic	For VSP 7400 Series: Hash Key = [Destination IP Address (32 bits), Source IP Address (32 bits), Source Port (8 bits), BackBone Destination MAC Address (48 bits), BackBone Source Mac Address (48 bits)] For all other platforms: Hash Key = [Source Port (8 bits), BackBone Destination MAC Address (48 bits), BackBone Source Mac Address (48 bits)]
Layer 2 Non-IP traffic	Hash Key = [Destination MAC Address (48 bits), Source MAC Address(48 bits)]

MultiLink trunking and autonegotiation interaction

To use MLT with the switch, you can have 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. After you use autonegotiation with MLT and not LACP, you need to ensure that all ports run at the same speed.

MLT configuration rules

Multilink trunks adhere to the following rules. Unless otherwise stated, these rules also apply to MLT with LACP.

- You cannot configure an MLT name that uses all numbers, for example, 222.
- 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.
- 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.

Multilink trunks have the following general features and requirements:

- Supports MLT groups with as many as 8 ports belonging to a single multilink trunk. For more information about the number of MLT groups supported for each hardware platform, see [Release Notes for VOSS](#).
- 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.

MLT with LACP LAG rules

The Link Aggregation Group (LAG) adheres 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.
- You can configure a LAG with up to 24 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 or VLANs manually without first disabling LACP.

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 switch 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 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 virtualized Switch Cluster through the SPBM cloud and are referred to as a Virtual Inter-Switch Trunk (vIST) 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 IPv4 Routing for VOSS](#).

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.

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 a Virtual Inter-Switch trunk (vIST), 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.

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, 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. 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 Virtual Inter-Switch Trunk (vIST)

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

*** Note:**

In the following figure, the pair of aggregation switches (E and F) are connected by a Virtual Inter-Switch Trunk (vIST). This means that although these two switches appear to be directly connected, they could be anywhere within the SPBM cloud.




Figure 2: SMLT aggregation switches and operations

You must connect SMLT aggregation switches through vIST. For example, user access switches B and C connect to the aggregation systems through multilink trunks split between the two aggregation systems. As shown above, the implementation of SMLT only requires two SMLT-capable aggregation switches.

Aggregation switches use vIST 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 vIST, use multiple links on vIST to ensure reliability and high availability. Use Gigabit Ethernet links for vIST connectivity to provide enough bandwidth for potential cross traffic.

For more information about vIST, see [Virtual interswitch trunk \(vIST\)](#) on page 13.

A vIST multilink trunk must contain at least two physical ports.

Other SMLT aggregation switch connections

The figure above includes end stations that connect to each of the switches. In this example, a, b1, b2, c1, c2, and d are clients and printers, while e and f are servers and routers.

User access switches B and C can use 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 vIST 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 a vIST 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 vIST as a remote SMLT, it is forwarded to the local port (the packet must not be sent to the vIST 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 switch obeys the MLT traffic distribution algorithm.

SMLT traffic flow examples

The following traffic flow examples are based on the figure above.

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 vIST 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 vIST because it is the shortest path. The link is treated as a standard link; SMLT and vIST 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 vIST if switch C sends it to E.

SMLT and vIST traffic flow example

In an SMLT environment, the two aggregation switches share the same forwarding database by exchanging forwarding entries using the vIST. The entry for 00:E0:7B:B3:04:00 is shown on switch C as an entry learned on MLT-1, but because SMLT Remote is true, this entry was actually learned

from switch 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 switch C destined for 00:E0:7B:B3:04:00, if the SMLT Remote status is true, the switch tries to send the packet to MLT-1 rather than through vIST. Traffic rarely traverses vIST 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

The switch fully supports IEEE 802.3ad LACP on MLTs and on a pair of SMLT systems.

With LACP the switch 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 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.

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:**

A VRRP virtual IP address must not be the same as the VLAN IP 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 vIST 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 vIST link to carry all the aggregated riser traffic. However, an enhancement to VRRP overcomes this issue by ensuring that the vIST 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 vIST 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, 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 vIST. This provides true load-sharing abilities.

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.
- 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. Always use SLPP in an SMLT environment.

*** Note:**

If SLPP is used in a vIST environment, it must be enabled on both the vIST peers. Because, when an SLPP packet of a vIST peer is looped through UNI ports to the other device, that device will shut down its UNI port due to receiving SLPP packets from its peer. A device's own SLPP packets will go over a vIST connection but will not be forwarded by its vIST peer back onto its UNI ports.

For square and full-mesh configurations that use a routed core, 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 switch does not support SLPP in an SMLT that uses LACP.

The SLPP design is to shut down the port where the SLPP packets originate or to shut down the vIST peer switch port, after the counter reaches the threshold. A loop can still occur after ports are shut down. SLPP can shut down 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, Spanning Tree, and NLB for VOSS](#).

SLPP Guard

Table 6: SLPP Guard product support

Feature	Product	Release introduced
For configuration details, see Configuring Link Aggregation, MLT, SMLT and vIST for VOSS .		
SLPP Guard	VSP 4450 Series	VOSS 6.1
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 6.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 6.1
	VSP 8400 Series	VOSS 6.1
	VSP 8600 Series	VSP 8600 6.2
	XA1400 Series	VOSS 8.0.50
		Not Supported on Flex-UNI

Use SLPP Guard with SMLT to provide additional loop protection to protect wiring closets from erroneous connections.

SLPP Guard requires Simple Loop Prevention Protocol (SLPP) to be configured in the core of the network. SLPP detects loops in the SMLT network. Because SMLT networks disable Spanning Tree (STP), Rapid Spanning Tree (RSTP), or Multiple Spanning Tree Protocol (MSTP) for participating

ports, SLPP Guard provides additional network loop protection, extending the loop detection to individual edge access ports.

SLPP Guard can be configured on MLT or LAG ports. If the edge switch with SLPP Guard enabled receives an SLPP-PDU packet on a port, SLPP Guard operationally disables the port for the configured timeout interval and appropriate log messages and SNMP traps are generated. If the disabled port does not receive any SLPP-PDU packets after the configured timeout interval expires, the port automatically reenables and generates a local log message, a syslog message, and SNMP traps, if configured.

MLT and SMLT configuration considerations

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

Note:

Static MAC is not supported against SMLT.

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.
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 the LACP for MLT procedures in this document.

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. Configure the LACP `smlt-sys-id` to be the base MAC address of one of the aggregate switches, and 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 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, set the multicast packets per-second value to 6000 pps.
- To avoid traffic loss when an SMLT goes down during a CP switchover, distribute multiple SMLT links across different slots.
- To avoid unnecessary processing, do not enable LACP on vISTs . Use VLACP if an optical network between the SMLT core switches requires a failure detection mechanism.

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 vIST 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 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. 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.

Note:

The switch 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 vIST 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 vIST 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

For the latest applicable scaling information, see [Release Notes for VOSS](#) for the version of the software running on the switch.

SMLT and VLACP configuration considerations

Use Virtual Link Aggregation Control Protocol (VLACP) for all SMLT access links configured as MultiLink Trunks to ensure both end devices can communicate. The switch does not support LACP and VLACP on the same links simultaneously.

VLACP for SMLT also protects against CPU failures by causing traffic to switch or reroute to the SMLT peer if the CPU fails or stops responding.

The following table provides the recommended values for VLACP in an SMLT environment:

Table 7: Recommended VLACP values

Parameter	Value
SMLT access	

Table continues...

Parameter	Value
Timeout	Short
Timer	500ms
Timeout scale	5
VLACP MAC	01:80:C2:00:00:0F
SMLT core	
Timeout	Short
Timer	500ms
Timeout scale	5
VLACP MAC	01:80:C2:00:00:0F
vIST	
Timeout	Long
Timer	10000
Timeout scale	3
VLACP MAC	01:80:C2:00:00:0F

SMLT with NNI ports configuration considerations

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

MLT and Private VLANs

The switch supports private VLANs and E-Tree configuration. MLT and private VLANs operate as follows:

- When using static MLT, configure the private-vlan as part of the overall MLT configuration, not at the port member level. All ports in the MLT use the private VLAN type for that MLT.
- When using LACP, configure the private-vlan at the interface level.
- When the private VLAN port type is trunk, the MLT automatically becomes tagged.
- If there are other non-private VLANs using the MLT configured as isolated, the following message is displayed: All non private VLANs using this interface will be removed once this mlt becomes isolated. Do you wish to continue (y/n) ?
- If there are other non-private VLANs using the MLT configured as promiscuous, the following message is displayed: All non private VLANs using this interface will be removed once this mlt becomes promiscuous. Do you wish to continue (y/n) ?

MLT and SMLT link aggregation configuration using the CLI

This section describes how to configure and manage link aggregation using the CLI, 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.

Configuring MLT

Perform this procedure to create and configure an MLT.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create an MLT:

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

3. **(Optional)** Set the private VLAN type for the MLT:

```
mlt <1-512> private-vlan <isolated|promiscuous|trunk>
```

4. Add a VLAN to the MLT:

```
mlt <1-512> vlan <1-4059>
```

5. **(Optional)** Change the name of the MLT:

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

6. Enable trunking on the MLT:

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

7. Enable the MLT:

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

8. Display the MLT configuration:

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

Example

Create MLT 40:

MultiLink Trunking and Split MultiLink Trunking

```
Switch:1(config)# mlt 40
```

Set the private VLAN type for MLT 40 to isolated:

```
Switch:1(config)# mlt 40 private-vlan isolated
```

Add VLAN 10 to the MLT:

```
Switch:1(config)# mlt 40 vlan 10
```

Enable the MLT:

```
Switch:1(config)# mlt 40 enable
```

Display the MLT configuration:

```
Switch:1(config)# show mlt 40
```

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

MLTID	IFINDEX	NAME	PORT TYPE	MLT ADMIN	MLT CURRENT	PORT MEMBERS	VLAN IDS
40	6183	MLT-40	access	norm	norm		10

MLTID	IFINDEX	DESIGNATED PORTS	LACP ADMIN	LACP OPER
40	6183	null	disable	down

MLTID	NAME	WHERE CREATED	LOCAL PORT MEMBERS	REMOTE PORT MEMBERS	WHICH PORTS PROGRAMMED IN DATA PATH
40	MLT-40	LOCAL			NONE

MLTID	IFINDEX	ENCAP DOT1Q	LOSSLESS	PVLAN	PVLAN TYPE	VID TYPE	FLEX-UNI
40	6183	disable	disable	enable	isolated	secondary	disable

Variable definitions

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

Variable	Value
<code>enable</code>	Creates and enables a new MLT.
<code>encapsulation dot1q</code>	Enables trunking on the MLT.
<code>member {slot/port[/sub-port] [-slot/port[/sub-port]] [...]}</code>	Adds ports to the MLT. Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized,

Table continues...

Variable	Value
	you must also specify the sub-port in the format slot/port/sub-port.
<1-512>	Specifies the MLT ID.
name WORD<0-20>	Configures the name for the MLT.
private-vlan {isolated promiscuous trunk}	Specifies a private VLAN type for the MLT.
vlan <1-4059>	Specifies a VLAN ID to add to the MLT.

Viewing MLT port members

View the port members in the specified MLT, or for all MLTs.

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. View information about the specified MLT:


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

Example

* Note:

The `show mlt [<1-512>]` command displays remote port members *only* if the MLTs are configured on vIST peers that use the same MLT ID.

If the MLT is between two devices that use the same MLT ID, but they are not vIST peers, `show mlt [<1-512>]` does not display the remote port members in the output.

In this first example, MLT 512 is an MLT between two vIST peers with the same ID used on both peers. In this case, the output displays the remote port members.

```
Switch:1(config)# show mlt 512
```

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

MLTID	IFINDEX	NAME	PORT TYPE	MLT ADMIN	MLT CURRENT	PORT MEMBERS	VLAN IDS
512	6655	MLT-512	trunk	norm	norm	1/24	1000 2000
MLTID	IFINDEX	DESIGNATED PORTS	LACP ADMIN	LACP OPER			
512	6655	null	disable	down			
		WHERE	LOCAL	REMOTE	WHICH PORTS PROGRAMMED		

MultiLink Trunking and Split MultiLink Trunking

MLTID	NAME	CREATED	PORT MEMBERS	PORT MEMBERS	IN DATA PATH
512	MLT-512	LOCAL	1/24	3/24	LOCAL
MLTID	IFINDEX	ENCAP DOT1Q	LOSSLESS		
512	6655	enable	disable		

In this second example, MLT 1 is an MLT between two devices that are *not vIST peers* so the output does not display the remote port members.

```
Switch:1(config)# show mlt 1
```

Mlt Info							
MLTID	IFINDEX	NAME	PORT TYPE	MLT ADMIN	MLT CURRENT	PORT MEMBERS	VLAN IDS
1	6144	MLT-1	trunk	norm	norm	1/5	1000 2000
MLTID	IFINDEX	DESIGNATED PORTS	LACP ADMIN	LACP OPER			
1	6144	1/5	disable	down			
MLTID	NAME	WHERE CREATED	LOCAL PORT MEMBERS	REMOTE PORT MEMBERS	WHICH PORTS PROGRAMMED IN DATA PATH		
1	MLT-1	LOCAL	1/5		LOCAL		
MLTID	IFINDEX	ENCAP DOT1Q	LOSSLESS				
1	6144	enable	disable				

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.

Adding ports to an MLT LAG

Perform this procedure to add ports to an MLT LAG.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Add ports to an MLT LAG:

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

Example

Add port 1/24 to MLT 40:

```
Switch:1(config)# mlt 40 member 1/24
```

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[/sub-port] [-slot/port[/sub-port]] [,...]}	Adds ports to the MLT. Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<1-512>	Specifies the MLT ID.
name WORD<0-20>	Configures the name for the MLT.
private-vlan {isolated promiscuous trunk}	Specifies a private VLAN type for the MLT.
vlan <1-4059>	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.

The shutdown port requires that you enter the interface GigabitEthernet configuration mode. However, to remove ports from an MLT LAG, you can enter any configuration mode.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

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

*** Note:**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable the ports:

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

3. Remove ports from an MLT LAG:

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

Example

Remove port 1/24 from MLT 40:

```
Switch:1(config)# interface GigabitEthernet 1/24
```

```
Switch:1(config-mlt)# shutdown port 1/24
```

```
Switch:1(config-mlt)# no mlt 40 member 1/24
```

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[/sub-port] [-slot/port[/sub-port]] [,...]}	Specifies the ports to add to the MLT. Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

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.

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
```

! Important:

If you want to remove SMLT configuration from the MLT on a vIST switch, first shut the SMLT MLT ports on its vIST peer.

Failure to perform this operation may lead to layer 2 loops.

! Important:

- If you are configuring SMLT with vIST, all the SMLT VLANs associated with vIST must have an I-SID.

Example

Create an SMLT on MLT 1:

```
Switch:1(config)#interface mlt 1
Switch:1(config-mlt)#smlt
```

Virtual interswitch trunk (vIST)

Creating a Virtual IST

Use this procedure to create a virtual IST (vIST).

vIST creates a virtualized channel through the SPBM cloud, and this channel connects two SMLT devices to form a virtualized Switch Cluster. Note that the SPBM cloud can consist of as few as two nodes.

! Important:

- When you create a vIST, VLANs assigned to SMLT ports must have an I-SID assigned.
- If you assign a VLAN to an I-SID on one SMLT-BEB node, then you must create the same VLAN and assign it to the same I-SID on the peer SMLT-BEB node even if no devices are connected to this second node.

- The vIST VLAN must also be associated with an I-SID. You can use the same vIST VLAN in another part of your network, but the I-SID associated with the vIST VLAN must not be used anywhere else.

*** Note:**

Simplified Virtual-IST (vIST) is for non-SPB customers who use SMLT with legacy IST. Simplified VIST is available for legacy multicast deployments only when the boot flag (`spbm-config-mode`) is disabled.

For more information, see [Configuring IP Multicast Routing Protocols for VOSS](#).

About this task

The following list provides an overview of the configuration steps:

- Enable SPBM and IS-IS globally.
- Configure SPBM and IS-IS.
- Configure an L2 VSN by assigning an I-SID to the C-VLAN, which is used by the vIST.

For information about SPBM and IS-IS, see [Configuring Fabric Basics and Layer 2 Services for VOSS](#).

! Important:

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of the header size increase when transmitting packets over the SPBM cloud.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create a vIST:

```
virtual-ist peer-ip <A.B.C.D> vlan <1-4059>
```

3. Verify your vIST configuration:

```
show virtual-ist
```

Example

The configuration example contains annotations (in parenthesis) to explain the configuration information.




Figure 4: SMLT with vIST configuration example

*** Note:**

This configuration example is for vIST peer B.

```
#
# ISIS SPBM CONFIGURATION
#

router isis
spbm 1
spbm 1 nick-name 0.00.16
spbm 1 b-vid 10,20 primary 10
spbm 1 smlt-virtual-bmac 00:00:84:04:01:01
spbm 1 smlt-peer-system-id a051.c6eb.c865
exit

#
# MLT CONFIGURATION
#

mlt 1 enable
mlt 1 member 1/28

mlt 2 enable
mlt 2 member 1/10
mlt 2 encapsulation dot1q

#
# VLAN CONFIGURATION
#

vlan members remove 1 1/28,1/10

(The rest of this VLAN section are all prerequisites for configuring a vIST.)
vlan create 10 type spbm-bvlan (This command and the next create the B-VLANs.)
vlan create 20 type spbm-bvlan
vlan create 33 type port-mstprstp 0
vlan i-sid 33 33
mlt 1 vlan 33

vlan create 100 type port-mstprstp 0 (Creates the C-VLAN for the vIST VLAN.)
vlan i-sid 100 100 (Creates a Layer 2 VSN for the vIST VLAN.)
```

MultiLink Trunking and Split MultiLink Trunking

```
interface Vlan 100

ip address 192.0.2.2 255.255.255.0 0  (Assigns an IP to the vIST VLAN.)
exit

#
# VIRTUAL vIST CONFIGURATION
#

virtual-ist peer-ip 192.0.2.1 vlan 100  (vIST configuration.)

#
# MLT INTERFACE CONFIGURATION
#

interface mlt 1
smlt
exit

interface mlt 2
isis
isis spbm 1
isis enable
exit

#
# PORT CONFIGURATION - PHASE II
#

interface GigabitEthernet 1/28
no shutdown
no spanning-tree mstp force-port-state enable
exit

interface GigabitEthernet 1/10
no shutdown
exit

#

# ISIS CONFIGURATION
#

router isis
is-type 11
manual-area 11.1111
exit
router isis enable
```

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 vIST VLAN on the other aggregation switch.
<1-4059>	Specifies the VLAN ID for this vIST.

Editing a virtual IST

If you need to change the virtual IST (vIST) `peer-ip` or `vlan`, use this procedure to delete the vIST first.

Note:

- You must disable IS-IS globally before deleting a vIST, and then re-enable it after creating a new vIST.
- You do not have to set the SMLT peer system ID or the virtual B-MAC to 0 before you change the vIST peer IP address or VLAN ID number.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Disable IS-IS globally:

```
no router isis enable
```

3. Delete the vIST:

```
no virtual-ist peer-ip
```

4. Create a vIST:

```
virtual-ist peer-ip <A.B.C.D> vlan <1-4059>
```

5. Enable IS-IS globally:

```
router isis enable
```

Configuring Simplified vIST in SMLT topologies

This procedure shows how to configure Simplified vIST in an SMLT environment. It includes steps to configure the following:

- Setting the boot config flag
- Configuring the vIST peer
- Enabling Simplified vIST

Important:

When you enable Simplified vIST with the `virtual-ist enable` command, two VLANs are automatically created to support vIST. The VLAN IDs are: 4086 and 4087.

Before you begin

SPBM must not be enabled on the vIST peers.

About this task

! Important:

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of a header size increase.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Disable the boot flag:

```
no boot config flags spbm-config-mode
```

The system responds with these messages:

```
Warning: Please save the configuration and reboot the switch for
this to take effect.
```

```
Warning: Please carefully save your configuration file before
rebooting the switch. Saving configuration file when spbm-config-
mode is changed to disable, removes SPBM configurations from the
configuration file.
```

3. Save the configuration, and then reboot the switch.

! Important:

Any change to the `spbm-config-mode` boot flag requires a reboot for the change to take effect.

4. Create the vIST VLAN:

```
vlan create <2-4059> type port-mstprstp <0-63>
interface vlan <1-4059>
ip address <A.B.C.D/X>
```

5. Configure the vIST peer address and VLAN:

```
virtual-ist peer-ip <A.B.C.D> vlan <1-4059>
```

6. Configure the SMLT MLT:

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

7. Configure the vIST MLT:

```

mlt <1-512> enable
mlt <1-512> member {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]}
mlt <1-512> encapsulation dot1q
interface mlt <1-512>
virtual-ist enable

```

*** Note:**

The **virtual-ist enable** command enables Simplified vIST and is only available when the **spbm-config-mode** boot flag is disabled.

8. Create a customer VLAN and assign the SMLT MLT ID:

```

vlan create <2-4059>
vlan mlt <1-4059> <1-512>
interface vlan <1-4059>
ip address <A.B.C.D/X>

```

Example

```

enable
configure terminal
no boot config flags spbm-config-mode

```

Save the configuration and reboot the switch.

```

virtual-ist peer-ip 192.0.2.1 vlan 50

mlt 3 enable
mlt 3 member 1/35,1/36
interface mlt 3
smlt
exit
mlt 5 enable
mlt 5 member 2/15,2/17
mlt 5 encapsulation dot1q
interface mlt 5
virtual-ist enable
exit
vlan create 50 type port-mstprstp 0
interface vlan 50
ip address 192.0.2.2 255.255.255.0 1
exit
vlan create 100
vlan mlt 100 3
interface vlan 100
ip address 198.51.100.1 255.255.255.0 2
exit

```

Viewing all ports configured for SMLT

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

Procedure

1. Enter Privileged EXEC mode:
enable
2. View all ports configured for SMLT:
show smlt mlt

Example

```
Switch:1#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.

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

1. Enter Privileged EXEC mode:
enable
2. View information about collision errors:
show mlt error collision [<1-512>]

Example

```
Switch:1#show mlt error collision 4
```

```
=====
                                     Mlt Collision Error
=====
MLT   -----COLLISIONS-----
ID    SINGLE  MULTIPLE LATE  EXCESSIVE
-----
40    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

1. Enter Privileged EXEC mode:

```
enable
```
2. View information about Ethernet errors:

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

Example

```
Switch:1#show mlt error main 40
```

```
=====
                                Mlt Ethernet Error
=====
MLT  ALIGN  FCS    IMAC   IMAC   CARRIER FRAMES  SQETEST DEFER
ID   ERROR   ERROR  TRNSMIT RECEIVE SENSE   TOOLONG ERROR  TRNSMSS
-----
40   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.

SLPP Guard configuration

This section provides the procedures to configure Simple Loop Prevention Protocol (SLPP) Guard.

! Important:

Enable SLPP Guard on the edge switches of an SMLT network and SLPP on the aggregation layer switches.

Configuring SLPP Guard

Configure SLPP Guard on MLT and LAG ports to provide additional network loop protection.

SLPP Guard is disabled by default.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

*** Note:**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable SLPP Guard on the port:

```
slpp-guard enable
```

3. **(Optional)** Configure the timeout value on the port:

```
slpp-guard timeout {<10-65535> | 0}
```

Example

```
Switch:1(config-if)#slpp-guard enable
Switch:1(config-if)#slpp-guard timeout 120
```

Variable definitions

Use the data in the following table to use the **slpp-guard** command.

Variable	Value
timeout <0 10-65535>	Specifies the time period, in seconds, for which SLPP Guard disables the port. After the timeout period expires, the switch reenables the port. The timeout value can be 0 or a value ranging from 10 to 65535. With a value of 0, the port remains disabled until it is manually re-enabled. The default timeout value is 60 seconds.

Reenabling an operationally disabled port

Reenable a port that has been operationally disabled by SLPP Guard.

*** Note:**

You cannot reenabling a disabled port if the timer count has not reached its timeout value. Either wait until it reaches the timeout or disable SLPP Guard for that port and then re-enable it.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

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

*** Note:**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Reenable SLPP Guard on the port:

```
slpp-guard enable
```

3. **(Optional)** Configure the timeout value on the port:

```
slpp-guard timeout {<10-65535> | 0}
```

Example

```
Switch:1(config-if)#slpp-guard enable
Switch:1(config-if)#slpp-guard timeout 120
```

Variable definitions

Use the data in the following table to use the `slpp-guard` command.

Variable	Value
timeout <0 10-65535>	Specifies the time period, in seconds, for which SLPP Guard disables the port. After the timeout period expires, the switch reenables the port. The timeout value can be 0 or a value ranging from 10 to 65535. With a value of 0, the port remains disabled until it is manually re-enabled. The default timeout value is 60 seconds.

Viewing SLPP Guard status

View current SLPP Guard settings.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display SLPP Guard status:

```
show slpp-guard {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]}
```

Example

```
Switch:1#>show slpp-guard

=====
                        SLPP Guard
=====
SLPP-guard Ethertype:  0x8102
=====
                        Port Interface
=====
Port      Link      Oper      SLPP-guard State  Timeout  TimerCount
-----
1/1       Up        Down     Disabled  N/A      60        N/A
1/2       Up        Down     Enabled   N/A      120       N/A
1/3       Up        Up       Enabled   Monitoring 60        N/A
1/4       Up        Down     Enabled   N/A      120       N/A
1/5       Down     Down     Disabled  N/A      60        N/A
1/6       Down     Down     Disabled  N/A      60        N/A
1/7       Down     Down     Disabled  N/A      60        N/A
1/8       Down     Down     Disabled  N/A      60        N/A
1/9       Down     Down     Disabled  N/A      60        N/A
1/10      Down     Down     Disabled  N/A      60        N/A
1/11      Down     Down     Disabled  N/A      60        N/A
1/12      Down     Down     Disabled  N/A      60        N/A
```

*** Note:**

The TimerCount column in the preceding example output indicates the time, in seconds, that elapses after SLPP Guard disables a port. When the value TimerCount value equals the Timeout value, the switch re-enables the port.

Variable definitions

Use the data in the following table to use the **show slpp-guard** command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

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.

Adding a multilink or LACP trunk

Perform this procedure to add a multilink or LACP trunk.

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. When you add the VLAN to the MLT, 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 tree, 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. In the **PrivateVlanType** box, select **trunk**, **isolated**, or **promiscuous**.
17. 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. 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. 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 Gbps, 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

Table continues...

Name	Description
RunningType	Specifies the MLT running type.
IfIndex	Specifies the interface of the trunk.
ClearLinkAggregate	Clears the link aggregate, disabling and reenabling the trunk.
DesignatedPort	Specifies the designated port of this trunk.
Aggregatable	Enables or disables link aggregation. 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.
PeerPortMembersList	Specifies the peer ports connected to the local ports of this trunk.
EntryOwner	Defines the owner of the MLT.
DatapathProgrammingState	Defines the datapath programming state of the MLT.
PrivateVlanType	Specifies the type of private VLAN for the MLT.
FlexUniEnable	Specifies whether the Flex UNI is enabled or disabled on the port. The default value is disable.

Adding ports to an MLT

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

Procedure

1. In the navigation tree, 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.

Creating a virtual IST using EDM

vIST creates a virtualized channel through the SPBM cloud, and this channel connects two SMLT devices to form a virtualized Switch Cluster. Note that the SPBM cloud can consist of as few as two nodes.

! Important:

- When you create a vIST, VLANs assigned to SMLT ports must have an I-SID assigned.
- If you assign a VLAN to an I-SID on one SMLT-BEB node, then you must create the same VLAN and assign it to the same I-SID on the peer SMLT-BEB node even if no devices are connected to this second node.
- The vIST VLAN must also be associated with an I-SID. You can use the same vIST VLAN in another part of your network, but the I-SID associated with the vIST VLAN must not be used anywhere else.

* Note:

Simplified Virtual-IST (vIST) is for non-SPB customers who use SMLT with legacy IST. Simplified VIST is available only for legacy multicast deployments when the boot configuration flag (`spbm-config-mode`) is disabled.

For more information, see [Configuring IP Multicast Routing Protocols for VOSS](#).

Before you begin

- Enable SPBM and IS-IS globally.

- Configure SPBM and IS-IS.
- Configure a Layer 2 VSN by assigning an I-SID to the C-VLAN, which is used by the vIST.

For information about SPBM and IS-IS, see [Configuring Fabric Basics and Layer 2 Services for VOSS](#).

About this task

! Important:

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of the header size increase when transmitting packets over the SPBM cloud.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN > MLT/LACP** folders.
2. Click the **Virtual IST** tab.
3. In the **PeerIp** field, type the peer IP address.
4. In the **VlanId** field, enter a VLAN ID.
5. Click **Apply**.

Virtual IST field descriptions

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

Name	Description
SessionStatus	Displays the status of the vIST session.
PeerIp	Specifies the peer IP address, which is the IP address of the vIST VLAN on the other aggregation switch.
VlanId	Configures a vIST VLAN ID number.

Editing a virtual IST

If you need to change the virtual IST **PeerIp** or **VlanId**, use this procedure to delete the vIST first.

* Note:

- You must disable IS-IS globally before deleting a vIST, and then re-enable it after creating a new vIST.
- You do not have to set the SMLT peer system ID or the virtual B-MAC to 0 before you change the virtual IST peer IP address or VLAN ID number.

Procedure

1. Disable IS-IS globally.
 - a. In the navigation pane, expand the following folders: **Configuration > IS-IS > IS-IS**.

- b. Click the **Globals** tab.
 - c. In the **AdminState** field, click **off**.
 - d. Click **Apply**.
2. Delete the vIST.
 - a. In the navigation pane, expand the following folders: **Configuration > VLAN > MLT/LACP**.
 - b. Click the **Virtual IST** tab.
 - c. Click **Apply**.
3. Create a vIST:
 - a. In the navigation pane, expand the following folders: **Configuration > VLAN > MLT/LACP**.
 - b. Click the **Virtual IST** tab.
 - c. In the **PeerIp** field, enter the peer IP address.
 - d. In the **VlanId** field, enter a VLAN ID.
 - e. Click **Apply**.
4. Enable IS-IS globally.
 - a. In the navigation pane, expand the following folders: **Configuration > IS-IS > IS-IS**.
 - b. Click the **Globals** tab.
 - c. In the **AdminState** field, click **on**.
 - d. Click **Apply**.

Configuring Simplified vIST in SMLT topologies

This procedure shows how to configure Simplified vIST in an SMLT environment. It includes steps to configure the following:

- Setting the boot config flag
- Configuring the vIST peer
- Enabling Simplified vIST

Important:

When you enable Simplified vIST with the `virtual-ist enable` command, two VLANs are automatically created to support vIST. The VLAN IDs are: 4086 and 4087.

Before you begin

SPBM must not be enabled on the vIST peers.

About this task

Important:

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of a header size increase.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Chassis** folders.
2. Click the **Boot Config** tab.
3. Clear the **EnableSpbmConfigMode** field to disable the boot flag.

The system responds with the following messages:

```
Warning: Please save the configuration and reboot the switch for
this to take effect.
```

```
Warning: Please carefully save your configuration file before
rebooting the switch. Saving configuration file when spbm-config-
mode is changed to disable, removes SPBM configurations from the
configuration file.
```

4. Click **Apply**.
5. Save the configuration, and then reboot the switch.

Important:

A change to the **EnableSpbmConfigMode** boot flag requires a reboot for the change to take effect.

6. Configure the *SMLT* MLT:
 - a. Expand the **Configuration > VLAN** folders.
 - b. Click **MLT/LACP**.
 - c. Click the **MultiLink/LACP Trunks** tab.
 - d. Click **Insert**.
 - e. In the **Id** box, type the ID number of the MLT.
 - f. In the **PortType** section, select access or trunk.
 - g. In the **Name** box, type a name for the MLT, or accept the default.
 - h. In the **PortMembers** box, click the (...) button.
 - i. In the **Port Editor: PortMembers** dialog box, select the desired ports.
 - j. Click **Ok**.
 - k. In the **VlanIdList** box, click the (...) button.
 - l. In the **VlanIdList** dialog box, select the desired VLANs.
 - m. Click **Ok**.

- n. In the **MltType** section, select splitMLT
- o. Click **Insert**.

The switch adds the SMLT MLT to the MultiLink/LACP Trunks tab in the MLT_LACP box.

7. Configure the *vIST* MLT:
 - a. Repeat steps [6.a](#) on page 89 to [6.o](#) on page 90 to configure the MLT.
 - b. Click **virtualistMLT** to enable Simplified vIST.
 - c. Click **Insert**.
8. Create the *vIST* VLAN:
 - a. Expand the **Configuration > VLAN > VLANs** folders.
 - b. In the **Basic** tab, click **Insert**.
 - c. In the **Id** box, enter an unused VLAN ID, or use the ID provided.
 - d. In the **MstpInstance** box, click the down arrow, and then choose an MSTI instance from the list.
 - e. In the **Type** box, select **byPort**.
 - f. In the **PortMembers** box, click the (...) button.
 - g. In the **Port Editor: PortMembers** dialog box, select the desired ports.
 - h. Click **OK**.
 - i. Click **Insert**.
 - j. Select the vIST VLAN from the list of VLANs, and then click **IP**.
 - k. Click **Insert**.
 - l. Configure the IP address for the vIST VLAN.
9. Repeat step [8](#) on page 90 to create an *SMLT* VLAN and assign the SMLT MLT ID to it. Do not use the vIST MLT ID.

SLPP Guard configuration

This section provide the procedures to configure Simple Loop Prevention Protocol (SLPP) Guard using EDM.

Important:

Enable SLPP Guard on the edge switches of an SMLT network and SLPP on the aggregation layer switches.

Configuring SLPP Guard

Configure SLPP Guard on MLT and LAG ports to provide additional network loop protection.

SLPP Guard is disabled by default.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
2. Click **SLPP**.
3. Click the **SLPP Guard** tab.
4. In the port row, double-click the cell in the **Enabled** column.
5. Select true from the drop-down list to enable SLPP Guard, or false to disable SLPP Guard for the port.
6. **(Optional)** In the port row, double-click the cell in the **Timeout** column.
7. **(Optional)** Type a value in the **Timeout** field.
8. Click **Apply**.
9. On the toolbar, click **Refresh** to update the work area data display.

SLPP Guard Ethernet type field descriptions

Use the data in the following table to use the SLPP Guard tab.

Name	Description
IfIndex	Specifies the port on which to configure SLPP Guard.
Enable	Enables (true) or disables (false) SLPP Guard for the port. The default is disabled.
Timeout	Specifies the time period, in seconds, for which SLPP Guard disables the port. After the timeout period expires, the switch re-enables the port. The timeout value can be 0 or a value ranging from 10 to 65535. With a value of 0, the port remains disabled until it is manually re-enabled. The default Timeout value is 60 seconds.
Status	Displays the SLPP Guard status for the port.
TimerCount	Indicates the time, in seconds, that elapses after SLPP Guard disables a port. When the TimerCount value equals the Timeout value, the switch re-enables the port.

Reenabling an operationally disabled port

Reenable a port that has been operationally disabled by SLPP Guard.

*** Note:**

You cannot reenabling a disabled port if the timer count has not reached its timeout value. Either wait until it reaches the timeout or disable SLPP Guard for that port and then re-enable it.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
2. Click **SLPP**.
3. Click the **SLPP Guard** tab.
4. In the port row, double-click the cell in the **Enabled** column.
5. Select true from the drop-down list to enable SLPP Guard, or false to disable SLPP Guard for the port.
6. **(Optional)** In the port row, double-click the cell in the **Timeout** column.
7. **(Optional)** Type a value in the **Timeout** field.
8. Click **Apply**.
9. On the toolbar, click **Refresh** to update the work area data display.

Viewing SLPP Guard status

View current SLPP Guard settings.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
2. Click **SLPP**.
3. Click the **SLPP Guard** tab.

SLPP Guard Ethernet type field descriptions

Use the data in the following table to use the SLPP Guard tab.

Name	Description
IfIndex	Specifies the port on which to configure SLPP Guard.
Enable	Enables (true) or disables (false) SLPP Guard for the port. The default is disabled.
Timeout	Specifies the time period, in seconds, for which SLPP Guard disables the port. After the timeout period expires, the switch re-enables the port. The timeout value can be 0 or a value ranging from 10 to 65535. With a value of 0, the port remains disabled until it is manually re-enabled.

Table continues...

Name	Description
	The default Timeout value is 60 seconds.
Status	Displays the SLPP Guard status for the port.
TimerCount	Indicates the time, in seconds, that elapses after SLPP Guard disables a port. When the TimerCount value equals the Timeout value, the switch re-enables the port.

MLT configuration examples

This chapter contains configuration examples for configuring MultiLink Trunking (MLT) and MLT with Link Aggregation Control Protocol (LACP) using the Command Line Interface (CLI).

* Note:

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

MultiLink Trunking

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

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




Figure 5: MLT within a VLAN

Switch A and B configuration

Configure MLT 10 on VLAN 100 on each device.

```
#
# MLT CONFIGURATION
#
mlt 10 enable
interface mlt 10
exit
#
# VLAN CONFIGURATION
#
vlan create 100 name VLAN-MLT-10 type port-mstprstp 0 color 1
vlan mlt 100 10
vlan members 100 1/1-1/4 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 6: MLT within a VLAN

The following procedures show you how to configure both switches.

Switch 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 0 color 1
vlan mlt 100 1
vlan mlt 100 10
vlan members 100 1/1-1/4 portmember
#
# PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 1/1
lACP key 10 aggregation enable
lACP enable
exit
interface GigabitEthernet 1/2
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.

*** Note:**

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Example 1: Switch-to-switch MLT

The following illustration shows two multilink trunks (MLT1 and MLT2) connecting three switches.




Figure 7: 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 8: 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 the switch. 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.

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




Figure 9: Client/server multilink trunk configuration

Chapter 5: Virtual Link Aggregation Control Protocol

Table 8: Virtual Link Aggregation Control Protocol (VLACP) product support

Feature	Product	Release introduced
For configuration details, see Configuring Link Aggregation, MLT, SMLT and vIST for VOSS .		
Virtual Link Aggregation Control Protocol (VLACP)	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50
VLACP Flap Detect and Damping	VSP 4450 Series	VOSS 7.1.5
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.1
	VSP 8200 Series	VOSS 8.1
	VSP 8400 Series	VOSS 8.1
	VSP 8600 Series	VSP 8600 8.0
	XA1400 Series	Not Supported

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.

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

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

For example, in [Figure 10: Problem description \(1 of 2\)](#) on page 99, 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 10: 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 11: 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.

The switch 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

VLACP considerations

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

- 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.
- If VLACP is enabled on a MACsec Key Agreement-enabled link, it takes approximately 30 seconds for the VLACP session to begin.

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).

VLACP Flap Detect and Damping

When there is instability or packet loss in a connection, the Virtual Link Aggregation Control Protocol (VLACP) state of the port flaps, causing services such as IP multicast to stop and start rapidly. This behavior causes system-wide instability, including high CPU utilization. In such cases, VLACP Flap Detect and Damping provides link flap detection capability, by configuring a specific time interval and frequency to count the number of VLACP flaps. After the system detects excessive VLACP flaps, it disables the specific VLACP port until the root cause is resolved.

For example, VLACP Flap Detect and Damping detects 3 VLACP flaps within 60 seconds, by default. After the system detects the first VLACP flap, the flap timer starts to count the number of

VLACP flaps occurring within 60 seconds. If the flap count reaches 3 before the timer ends, the system disables the specific VLACP port, and generates a Simple Network Management Protocol (SNMP) trap. For more information about SNMP traps, see [Troubleshooting VOSS](#).

If you view the port state information, the port displays as operationally down and the reason provided is VLACP_FLAP. The following example shows a port that is disabled because of VLACP flapping; only the relevant sections of the show command output is included.

```
Switch:1#show interface gigabitethernet 1/1
=====
                        Port Interface
=====
PORT      INDEX  DESCRIPTION      LINK  PORT  PHYSICAL      STATUS
NUM       INDEX  DESCRIPTION      TRAP  LOCK  ADDRESS        ADMIN  OPERATE
-----
1/1       192    1000BaseTX       true  false  e4:5d:52:3c:64:00  up    down
=====
                        Port State
=====
PORT NUM  ADMINSTATUS  PORTSTATE  REASON          DATE
-----
1/1      up           down       VLACP_FLAP    09/19/19 02:22:58
=====
```

*** Note:**

VLACP Flap Detect and Damping does not support auto-recovery of VLACP ports; you must reenble the ports manually.

VLACP Flap Detect and Damping Considerations

Use the following information when configuring VLACP Flap Detect and Damping.

- Do not enable VLACP Flap Detect and Damping on Link Aggregation Control Protocol (LACP) enabled ports.

*** Note:**

Link flap detection takes priority if you enable both VLACP Flap Detect and Damping and link flap detection on the same port.

VLACP Configuration using CLI

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

Configure 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.

! Important:

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

About this task

Use the following information to prevent flooding VLACP packets across a defaulted switch:

- Use the default MAC address, 01:80:c2:00:11:00, for end-to-end connections that traverse an intermediate network.
- Use the reserved multicast MAC address 01-80-c2-00-00-0f for directly-connected, peer-to-peer links.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

* Note:

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-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.

3. Enable VLACP on a port:

```
vlacp enable
```

Example

Configure VLACP on port 1/1:

```
Switch:1# configure terminal
Switch:1# interface GigabitEthernet 1/2
Switch:1# vlapc fast-periodic-time 400 timeout short
Switch:1# vlapc 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.
ethertype <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.
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. The default is 01:80:c2:00:11:00.
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.

View the VLACP Port Configuration

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

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View the VLACP port configuration for all interfaces:

```
show vlacp interface gigabitethernet [vid <1-4059>] [{slot/port[/sub-port]} [-slot/port[/sub-port]] [,...]]
```

Example

```
Switch:1>show vlacp interface gigabitethernet
=====
                        VLACP Information
=====
INDEX  ADMIN  OPER   PORT   FAST  SLOW   TIMEOUT TIMEOUT ETHER   MAC
      ENABLED  ENABLED STATE  TIME  TIME   TIME    SCALE  TYPE    ADDR
-----
1/1    false  false  DOWN   200   30000 long    3      0x8103 01:80:c2:00:11:00
1/2    false  false  DOWN   200   30000 long    3      0x8103 01:80:c2:00:11:00
1/3    false  false  DOWN   200   30000 long    3      0x8103 01:80:c2:00:11:00
1/4    false  false  DOWN   200   30000 long    3      0x8103 01:80:c2:00:11:00
1/5    false  false  DOWN   200   30000 long    3      0x8103 01:80:c2:00:11:00
--More-- (q = quit)
```

```
=====
                        VLACP Flap Detect Information
=====
INDEX  FLAP   FLAP   FLAP   TOTAL   FIRST-FLAP   LAST-FLAP
      DETECT  FREQ  INTERVAL  FLAP    TIME          TIME
-----
1/1    false  3      60       0       -- --        -- --
1/2    false  3      60       0       -- --        -- --
1/3    false  3      60       0       -- --        -- --
1/4    false  3      60       0       -- --        -- --
1/5    false  3      60       0       -- --        -- --
--More-- (q = quit)
```

```
Switch:1>show vlacp interface gigabitethernet 2/11 vid 2
=====
                        VLACP Information
=====
INDEX  ADMIN  OPER   PORT   FAST  SLOW   TIMEOUT TIMEOUT ETHER   MAC
      ENABLED  ENABLED STATE  TIME  TIME   TIME    SCALE  TYPE    ADDR
-----
2/11   true   true   up     500   30000 short   3      0x8103 01:80:c2:00:00:0f
```

```
=====
                        VLACP Flap Detect Information
=====
INDEX  FLAP   FLAP   FLAP   TOTAL   FIRST-FLAP   LAST-FLAP
      DETECT  FREQ  INTERVAL  FLAP    TIME          TIME
-----
2/11   true   10     10       4     09/20/11 08:24:10  09/20/11 08:26:10
```

```
Switch:1>show vlacp interface gigabitethernet 1/7
=====
                        VLACP Information
=====
INDEX  ADMIN  OPER   PORT   FAST  SLOW   TIMEOUT TIMEOUT ETHER   MAC
      ENABLED  ENABLED STATE  TIME  TIME   TIME    SCALE  TYPE    ADDR
-----
1/7    false  false  DOWN   200   30000 long    3      0x8103 01:80:c2:00:11:00
```

```
=====
                        VLACP Flap Detect Information
=====
INDEX  FLAP   FLAP   FLAP   TOTAL   FIRST-FLAP   LAST-FLAP
      DETECT  FREQ  INTERVAL  FLAP    TIME          TIME
-----
1/7    false  3      60       0       -- --        -- --
```


Variable definitions

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

Variable	Value
vid <1-4059>	<p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p> <p>The VLAN ID is in one of the following formats: A single VLAN ID (vlan-id), a range of VLAN IDs [(vlan-id)-(vlan-id)] or a series of VLAN IDs (vlan-id, vlan-id, vlan-id).</p>
{slot/port[/sub-port] [-slot/port[/sub-port]] [...]}	<p>Specifies a port or list of ports to show only the VLACP information for those ports.</p> <p>Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

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. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable VLACP globally:

```
vlacp enable
```

3. Disable VLACP globally:

```
no vlacp enable
```

Example

Enable VLACP globally:

```
Switch:1(config)# vlacp enable
```

Configure VLACP Flap Detect and Damping on a Port

About this task

Perform the following procedure to control link state changes on VLACP ports. By default, VLACP Flap Detect and Damping is disabled on all VLACP ports.

! Important:

Do not enable VLACP Flap Detect and Damping on Link Aggregation Control Protocol (LACP) enabled ports.

Before you begin

- To configure VLACP Flap Detect and Damping on sub-ports, enable channelization. For more information about channelization, see [Administering VOSS](#).
- As a best practice, to modify an existing VLACP Flap Detect and Damping configuration, first disable VLACP on the interface.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

*** Note:**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable VLACP Flap Detect and Damping on the specified port:

```
vlacp flap-detect enable
```

3. Configure the time interval (in seconds) to record VLACP flaps:

```
vlacp flap-interval <10-600>
```

4. Configure the VLACP flap frequency:

```
vlacp flap-frequency <3-30>
```

5. View VLACP Flap Detect and Damping configuration on the specific port:

```
show vlacp interface gigabitethernet {slot/port[/sub-port]}[-slot/
port[/sub-port]][,...]
```

Example

Configure VLACP Flap Detect and Damping on port 1/3:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface GigabitEthernet 1/3
Switch:1(config-if)#vlacp flap-detect enable
Switch:1(config-if)#vlacp flap-interval 120
Switch:1(config-if)#vlacp flap-frequency 6
Switch:1(config-if)#show vlacp interface gigabitethernet 1/3
=====
                                VLACP Information
=====
INDEX  ADMIN   OPER    PORT   FAST    SLOW    TIMEOUT  TIMEOUT  ETHER   MAC
      ENABLED  ENABLED STATE  TIME    TIME    TIME     SCALE    TYPE   ADDR
-----
1/3   false   false   DOWN   200     30000   long     3        0x8103  01:80:c2:00:11:00
=====
                                VLACP Flap Detect Information
=====
INDEX  FLAP    FLAP    FLAP    TOTAL    FIRST-FLAP    LAST-FLAP
      DETECT  FREQ   INTERVAL  FLAP    TIME          TIME
-----
1/3   enable   6      120      6       01/03/19 15:11:52  01/03/19 15:12:12
=====
```

Variable Definitions

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

Variable	Value
flap-detect enable	Enables VLACP Flap Detect and Damping on the specified port. By default, VLACP Flap Detect and Damping is disabled on all ports.
flap-interval<10-600>	Configures VLACP Flap Detect and Damping time interval in seconds. The default value is 60 seconds.
flap-frequency<3-30>	Configures the VLACP Flap count permitted during the configured time interval. The default value is 3.

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 tree, 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.

Configure VLACP on a port

Perform the following procedure to enable VLACP on a port. VLACP periodically checks the end-to-end condition of a point-to-point connection. You can also configure VLACP Flap Detect and Damping to control link state changes on VLACP ports. By default, VLACP Flap Detect and Damping is disabled on all VLACP ports.

Important:

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

Before you begin

- To configure VLACP Flap Detect and Damping on sub-ports, enable channelization. For more information about channelization, see [Administering VOSS](#).
- As a best practice, to modify an existing VLACP or VLACP Flap Detect and Damping configuration, first disable VLACP on the interface.

About this task

Use the following information to prevent flooding VLACP packets across a defaulted switch:

- Use the default MAC address, 01:80:c2:00:11:00, for end-to-end connections that traverse an intermediate network.
- Use the reserved multicast MAC address 01-80-c2-00-00-0f for directly-connected, peer-to-peer links.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **VLACP** tab.
5. Configure the parameters, as required.
6. Select **Apply**.

VLACP Field Descriptions

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

Name	Description
AdminEnable	Enables VLACP for the port. The default is disabled.
OperEnable	Specifies the VLACP operational status for the port. The default is false.
FastPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using short timeouts. Use the same value for all LACP-enabled ports. The default value is 200.
SlowPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using long timeouts. Use the same value for all LACP-enabled ports. 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 from the periodic time for all VLACP-enabled ports. $\text{Timeout} = \text{PeriodicTime} \times \text{TimeoutScale}$. The default value is 3.
EtherType	Specifies the VLACP identification. The ID is in hexadecimal format. The default value is 0x8103.
EtherMacAddress	Specifies the multicast MAC address exclusively used for VLACPDUs. The default is 01:80:c2:00:11:00.
PortState	Specifies the VLACP port state.

Table continues...

Name	Description
FlapDetectEnable	Enables VLACP Flap Detect and Damping on the port. By default, VLACP Flap Detect and Damping is disabled on all ports.
FlapFrequency	Specifies the number of VLACP flaps detected on the port. The port is shut down if the recorded flaps are within the configured time interval. The default value is 3.
FlapInterval	Specifies the time interval (in seconds) during which the VLACP flaps on the port are monitored. The default value is 60 seconds.
TotalFlapCount	Shows the total number of VLACP flaps detected on the port since you enabled VLACP Flap Detect and Damping or cleared the counters and timers.
FirstFlapTimeStamp	Shows the timestamp of the first VLACP flap detected on the port since you enabled VLACP Flap Detect and Damping or cleared the counters and timers.
LastFlapTimeStamp	Shows the timestamp of the latest VLACP flap detected on the port since you enabled VLACP Flap Detect and Damping or cleared the counters and timers.
FlapClearStats	Clears all VLACP Flap Detect and Damping counters and timers.

Configuring VLACP on an Extreme Integrated Application Hosting Port

About this task

Perform this procedure to enable Virtual Link Aggregation Control Protocol (VLACP) on an Extreme Integrated Application Hosting (IAH) port. VLACP periodically checks the end-to-end health of a point-to-point connection.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Insight Port**.
2. Select the IAH port you want to configure.
3. Select the **VLACP** tab.
4. Select **AdminEnable**.
5. Configure the other parameters as required.
6. Click **Apply**.

VLACP Field Descriptions

Use data in the following table to configure the VLACP tab.

Name	Description
AdminEnable	Enables VLACP on the Extreme Integrated Application Hosting (IAH) port. The default is disabled.
OperEnable	Specifies the VLACP operational status for the IAH port. The default is false.
FastPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using short timeouts. Set the same value for all LACP enabled IAH ports. The default value is 200.
SlowPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using long timeouts. Set the same value for all LACP enabled IAH ports. The default value is 30000.
Timeout	Specifies the timeout control value. The default value is long.
TimeoutScale	Specifies the value used to calculate timeout duration from the periodic time for all VLACP enabled IAH ports. Timeout = PeriodicTime x TimeoutScale. 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	Specifies the VLACP port state.

Chapter 6: Link-state tracking (LST)

Table 9: Link-state tracking (LST) product support

Feature	Product	Release introduced
For configuration details, see Configuring Link Aggregation, MLT, SMLT and vIST for VOSS .		
Link-state tracking (LST)	VSP 4450 Series	VOSS 7.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 7.0
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 7.0
	VSP 8400 Series	VOSS 7.0
	VSP 8600 Series	VSP 8600 8.0 demonstration feature
	XA1400 Series	Not Supported

*** Note:**

DEMO FEATURE - Link-state tracking (LST) is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

Link-state tracking (LST) Overview

Link-state tracking (LST) binds the link state of multiple interfaces, creating LST groups with upstream (to-be-followed) and downstream (to-follow) interfaces. LST monitors the state of upstream interfaces and automatically transfers the upstream state to the downstream interfaces. If all the upstream interfaces in a LST group are down, the downstream interfaces are administratively configured as down after approximately five seconds. If any upstream interface in a LST group is up, the downstream interfaces are not affected. The role of the LST group is to keep the downstream interfaces in the same state as the upstream interface.

An interface can be an aggregation of ports, multi-link trunks (MLT) or link aggregation groups (LAG). Interfaces can only belong to one LST group. You can configure LST using CLI or EDM. LST receives updates from Port Manager, MLT, and VLACP regarding the upstream state of ports and trunks in the group.

LST can detect a link failure of upstream interfaces and shutdown downstream interfaces, eliminating loss of traffic and allowing the source to reroute traffic. When a LST group disables a downstream interface, the interface can only be enabled by the LST group. You can recover the downstream interfaces that LST disabled by removing the interfaces from the LST group or by disabling the LST group. You can administratively enable or disable LST group downstream interfaces with shutdown commands. A LST group cannot enable ports that you administratively disabled.

*** Note:**

If you administratively enable an interface which LST disabled, only the administrative status of the interface changes. The interface remains disabled until the LST group enables the interface, or until you remove the interface from the LST group.

MLT interactions

For MLTs, a last-link-down event triggers a down operational state and a first-link-up event triggers an up operational state. You cannot delete a MLT that is a member of a LST group.

LAG and LACP interactions

For LAGs, a static LAG trunk ID associated with a LACP administrative key can be a member of a LST group. You can add LAGs to a LST group by specifying the LAG trunk ID. You cannot break the association between trunk ID, LACP key, and ports while the LAG trunk is in a LST group. For LACP, you cannot add ports with link-aggregation enabled to a LST group, or enable link-aggregation on interfaces already in a LST group.

VLACP interactions

You can add LST group interfaces configured with VLACP. For upstream interfaces with VLACP enabled, when the physical link is up with a VLACP partner, the operational state is up. Otherwise the operational state is down. If VLACP is enabled, the value of the VLACP have partner field and the link status correspond to the operational state of upstream interfaces. For upstream interfaces with VLACP disabled, the up and down operational states correspond directly with the physical link.

SLPP Guard interactions

You can add LST group interfaces configured with SLPP Guard. When LST disables a port that is already disabled by SLPP Guard, the interface is unblocked by SLPP Guard and the blocking timer clears.

BPDU Guard and MACsec interactions

You can add LST group interfaces configured with BPDU Guard or MACsec. BPDU Guard and MACsec can enable or disable ports administratively. An interface is enabled if both LST and BPDU Guard or MACsec consider the port enabled. If BPDU Guard or MACsec disables the port, the port remains down and does not link up.

LST configuration using CLI

*** Note:**

DEMO FEATURE - Link-state tracking (LST) is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

Configure Link-state tracking (LST) groups composed of upstream and downstream interfaces. Valid LST group members are switch ports, multi-link trunks (MLT), or link aggregation groups (LAG).

Configuring LST

*** Note:**

DEMO FEATURE - Link-state tracking (LST) is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

Use this procedure to configure LST groups.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the LST group upstream and downstream interface members:

```
link-state group <1-48> <upstream | downstream> interface
gigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]}
```

3. Configure the LST group upstream and downstream mlt members:

```
link-state group <1-48> <upstream | downstream> mlt <1-512>
```

4. Enable the group:

```
link-state group <1-48> enable
```

5. Display the status of the groups:

```
show link-state group <1-48> [detail]
```

Example

The following example shows a LST group created with ports 1–4 and mlt 1 as upstream members and ports 11–14 and mlt 2 as downstream members. The LST group is enabled. Because port 1 is up, the LST group receives an up operational state and the downstream interfaces are enabled.

Upstream ports 2, 4 and mlt 1 have VLACP admin enabled and are listed in the VLACP Upstream State section.

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# link-state group 1 upstream interface gigabitEthernet 1/1-1/4
Switch:1(config)# link-state group 1 upstream mlt 1
Switch:1(config)# link-state group 1 downstream interface gigabitEthernet 1/11-1/14
Switch:1(config)# link-state group 1 downstream mlt 2
Switch:1(config)# link-state group 1 enable
Switch:1(config)# show link-state group 1
=====
Link State Tracking General Info
=====
Group                1
Status               Enabled
Operational Status   UP
-----

Switch:1(config)# show link-state detail
=====
Link State Tracking Detailed Info
=====
Group:                1
Status:               Enabled
Operational Status:   UP
VLACP Upstream State: Active on ports: 1/2, 1/4
                    Active on Trunks: 1
Upstream Ports:
1/1 (UP)  1/2 (DW)  1/3 (DW)  1/4 (DW)
Upstream Trunks:
1 (UP)
Downstream Ports:
1/11 (UP) 1/12 (DW) 1/13 (DW) 1/14 (UP)
Downstream Trunks:
2 (DW)
-----
Group                : 2
Status               : Disabled
Operational Status   : N/A
VLACP Upstream State: : N/A
Upstream Ports:      : Not Configured
Upstream Trunks:     : Not Configured
Downstream Ports:    : Not Configured
Downstream Trunks:   : Not Configured
-----
```

Variable definitions

Use the data in the following table to use the `link-state` command.

Variable	Value
enable	Activates the action specified for the LST group or specified interfaces.
group <1-48>	Specifies a LST group ID.
interface gigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [...]}	Adds ports to the specified upstream or downstream LST group.

Table continues...

Variable	Value
	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
mlt <1–512>	Adds a MLT to the specified upstream or downstream LSTgroup.
<upstream downstream>	Specifies the upstream or downstream interfaces in the LST group.

Use the data in the following table to use the **show link-state** command.

Variable	Value
detail	Displays the specified LST group status as enabled or disabled and if the operational status is up or down. Detail displays LST group additional information.

LST configuration using EDM

*** Note:**

DEMO FEATURE - Link-state tracking (LST) is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

Configure Link-state tracking (LST) groups composed of upstream and downstream interfaces. Valid LST group members are switch ports, multi-link trunks (MLT), or link aggregation groups (LAG).

Configuring LST

*** Note:**

DEMO FEATURE - Link-state tracking (LST) is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

Use this procedure to configure LST groups.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **Link State Tracking**
3. Click the **Link State Tracking** tab.
4. Choose a **GroupId** row to configure.
5. To add upstream ports to the LST group, double click the **UpstreamPortList** field, select the upstream ports to add, then click **OK**.
6. To add downstream ports to the LST group, double click the **DownstreamPortList** field, select the upstream ports to add, then click **OK**.
7. To add an upstream MLT to the LST group, double click the **UpstreamMltList** field and enter an upstream MLT ID.
8. To add a downstream MLT to the LST group, double click the **DownstreamMltList** field and enter an upstream MLT ID.
9. To enable a LST group, double click the **Enabled** field and select **true**.
10. To activate the LST group configuration, click **Apply**.
11. View the **OperState** field of the LST groups to verify the current operating state.

Link State Tracking field descriptions

Use the data in the following table to use the Link State Tracking tab.

Name	Description
GroupId	Specifies the LST group number between 1 and 48.
Enabled	Specifies if the LST group is enabled. Values are true or false. Default is false.
UpstreamPortList	Specifies upstream interface ports in the LST group.
DownstreamPortList	Specifies downstream interface ports in the LST group.
UpstreamMltList	Specifies an upstream multi-link trunk in the LST group.
DownstreamMltList	Specifies a downstream multi-link trunk in the LST group.
OperState	Displays the current operating status of the LST group. Values are up, down, or notConfigured.

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.
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 network handshaking protocol that provides a means to aggregate multiple links between appropriately configured devices.
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.
SFP	A hot pluggable, small form-factor pluggable (SFP) transceiver, which is used in Ethernet applications up to 1 Gbps.
Simple Loop Prevention Protocol (SLPP)	Simple Hello Protocol that prevents loops in a Layer 2 network (VLAN).
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 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 IST (vIST)	Virtual Inter-Switch Trunk (vIST) improves Layer 2 and Layer 3 resiliency by using a virtualized IST channel through the SPBM cloud. The vIST channel is always up as long as there is SPBM connectivity between the vIST peers.
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.