AVAYA

Configuration — Link Aggregation, MLT, and SMLT
Avaya Virtual Services Platform 9000

Contents

Chapter 1: New in this release

The following section details what's new in *Avaya Virtual Services Platform 9000 Configuration — Link Aggregation, MLT, and SMLT, NN46250-503* for Release 3.0:

- Features on page 5
- Other changes on page 5

Features

There are no feature changes.

Other changes

See the following sections for information about changes that are not feature-related.

SMLT and interswitch trunking

SMLT and interswitch trunking on page 19 is updated with information about MAC address and Address Resolution Protocol (ARP) entries that an SMLT interfaces learns.

SMLT, MSTP, and RSTP

SMLT, MSTP, and RSTP on page 18 is updated to reflect traffic loss during Spanning Tree Protocol (STP) convergence.

SMLT and SLPP

SMLT and SLPP on page 25 is updated to identify the potential isolation of an edge switch.

Chapter 2: Introduction

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

For information about terms and acronyms used in this document, see *Avaya Virtual Services Platform 9000 Terminology* (NN46250-101).

Chapter 3: Link aggregation fundamentals

This chapter describes link aggregation concepts and the features supported on the Avaya Virtual Services Platform 9000.

Link aggregation overview

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

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

- MultiLink Trunking (MLT)—a statically configured link bundling method. MLT is not standards based, but it interoperates with static link methods of other vendors. If you aggregate links from multiple modules together, MLT can provide not only port redundancy, but also module redundancy.

- IEEE 802.3ad based link aggregation, through the Link Aggregation Control Protocol (LACP), dynamically aggregates links as they become available to a trunk group. Link Aggregation Control Protocol dynamically detects whether links can be aggregated into a link aggregation group (LAG) and does so after links become available. Link Aggregation Control Protocol also provides link integrity checking at Layer 2 for all links within the LAG.

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

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

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

These SMLT devices form a Switch Cluster and are referred to as an interswitch Trunk (IST) Core Switch pair.

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

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

MultiLink Trunking

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

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

- MLT traffic distribution algorithm on page 10
- Multicast flow distribution over MLT on page 11
- MLT and MLT with LACP configuration rules on page 13
- LAG rules on page 14
- MLT network topology and configuration reference on page 71

MLT traffic distribution algorithm

You can use a multilink trunk to aggregate bandwidth between two switches. The Virtual Services Platform 9000 uses one of two algorithms to determine which active port in the multilink trunk to use for each packet. The MLT algorithms provide load sharing, not load balancing. The MLT algorithms ensure that each packet in a flow does not arrive out of sequence, and that a flow always traverses the same link path.

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

For more information about port and algorithm calculations, see Viewing the MLT port calculated by the MLT hash algorithm on page 61.

The Virtual Services Platform 9000 uses an enhanced traffic distribution algorithm, which ensures proper traffic distribution in all customer networks. The system determines MultiLink Trunking hashing for a module type at the ingress port, although the effect is seen at the egress MLT. That is, the traffic ingressing a port determines the destination is an MLT and sends the traffic to the correctly hashed MLT port, regardless of MLT port type or even a mixed port type.

Therefore, in a mixed chassis, the algorithm used is based on the ingress port for the traffic, not the MLT configuration. The new hash algorithm for IPv4 TCP/UDP traffic is as follows

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

For non-TCP/UDP IPv4 traffic

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

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

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

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

MultiLink trunking and autonegotiation interaction

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

Multicast flow distribution over MLT

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

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

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

After you change the state parameters of multicast flow distribution, minor traffic interruptions can occur.

For more information about the commands used to configure multicast flow distribution over MLT, see *Avaya Virtual Services Platform 9000 Configuration — IP Multicast Routing Protocols, NN46250-504*.

Multicast distribution algorithm

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

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

The following is an example:

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

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

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

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

The port is calculated by using the following:

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

Algorithm example

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

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

X.Y.Z.1

X.Y.Z.2

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

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

Multicast traffic distribution

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

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

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

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

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

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

For more information about the commands used to configure multicast flow distribution over MLT, see *Avaya Virtual Services Platform 9000 Configuration — IP Multicast Routing Protocols, NN46250-504*.

MLT and MLT with LACP configuration rules

Virtual Services Platform 9000 multilink trunks adhere to the following rules. The rules also apply to MLT with LACP.

- MLT is supported on 10Base-T, 100Base-TX, Gigabit Ethernet, and 10 Gigabit Ethernet module ports.

- Multilink trunk ports support mixed speed links, for example, using a 1Gb SFP in a 10Gb port. 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.

- The media type of MLT ports can be different; a mix of copper and fiber are allowed.

- All multilink trunk ports must be in the same Spanning Tree Group (STG) unless the port is tagged. Use tagging so ports can belong to multiple STGs, as well as multiple VLANs.

- After the port is made a member of MLT, it inherits the properties of the MLT and hence the STG properties are inherited from the VLAN associated with that MLT. After you remove the port from MLT or after you delete the MLT, the ports are removed from the MLT STG and added into the default STG.

- MLT is compatible with Multiple Spanning Tree Protocol (MSTP) (IEEE 802.1s) and Rapid Spanning Tree Protocol (RSTP) (IEEE 802.1w).

- Tagging (IEEE 802.1Q) is supported on a multilink trunk.

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

- Supports up to 512 MLT groups with as many as 16 ports belonging to a single multilink trunk.

- MLT ports can span modules, providing module redundancy.

- Apply filters individually to each port in a multilink trunk.

With MSTP or RSTP enabled, ports in the same multilink trunk operate as follows:

- The designated port sends the Bridge Protocol Data Unit (BPDU).

- The multilink trunk port ID is the ID of the lowest numbered port.

- If identical BPDUs are received on all ports, the multilink trunk mode is forwarding.

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

LAG rules

The Virtual Services Platform 9000 Link Aggregation Groups (LAG) adhere to the following rules:

- All LAG ports operate in full-duplex mode.
- All LAG ports operate at the same data rate.
- All LAG ports must belong to the same set of VLANs.
- Link aggregation is compatible with MSTP, and RSTP.
- Assign all LAG ports to the same MSTP or RSTP groups.
- Ports in a LAG can exist on different modules.
- Support a maximum of 32 link aggregation groups.
- You can configure a LAG with up to 16 ports, but only a maximum of 8 can be active at a time.
- After you configure a multilink trunk with LACP, you cannot add or delete ports or VLANs manually without first disabling LACP.

MultiLink Trunking with LACP

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

IEEE 802.3ad overview

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

The following illustration shows the major functions of IEEE 802.3ad defined as multiple link aggregation.
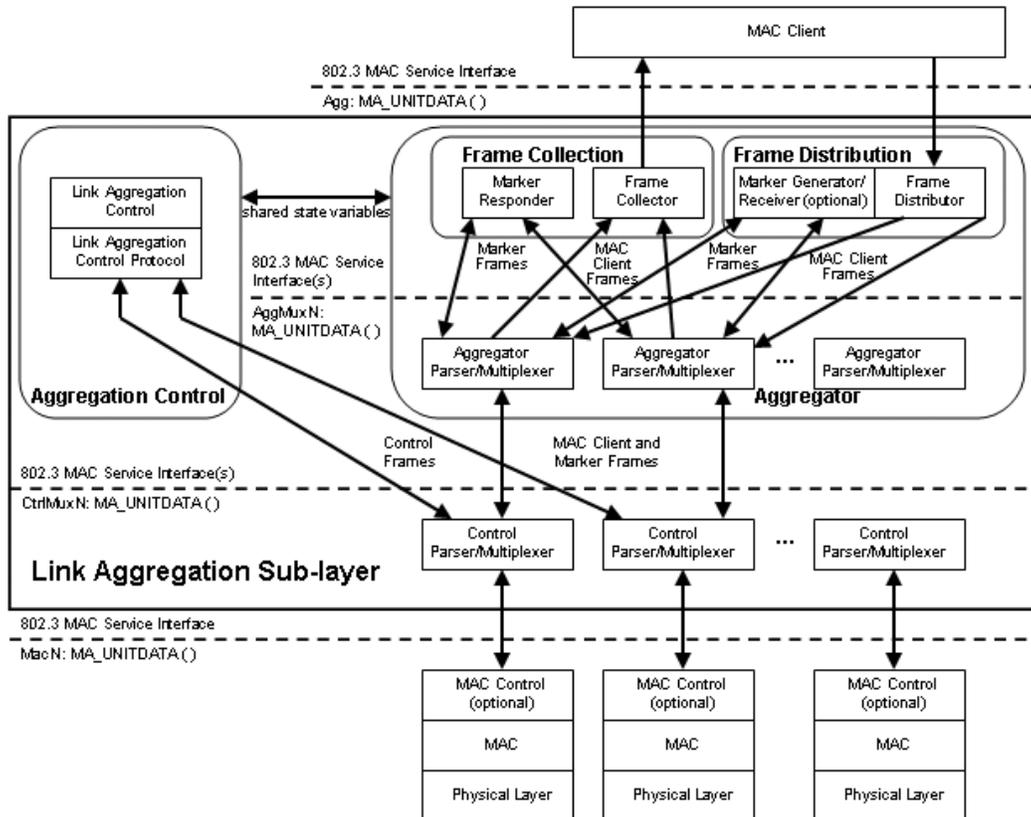
Figure 1: Link aggregation sublayer (according to IEEE 802.3ad)

The link aggregation sublayer comprises the following functions:

- frame distribution

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

  Frame distribution also includes an optional Marker Generator/Receiver used for the Marker Protocol. Virtual Services Platform 9000 only implements the Marker Receiver function. For more information about the frame distribution algorithm, see MLT traffic distribution algorithm on page 10.

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

802.3ad link aggregation principles

Use link aggregation to group ports together to form a link group to another device. Link groups increase aggregate throughput between devices and provide link redundancy.

Link aggregation employs the following principles and concepts:

- A MAC client communicates with a set of ports through an aggregator, which presents a standard IEEE 802.3 service interface to the MAC client. The aggregator binds to one or more ports within a system.

- The aggregator distributes frame transmissions from the MAC client to various ports, collects received frames from the ports, and transparently passes the frames to the MAC client.

- A system can contain multiple aggregators serving multiple MAC clients. A port binds to a single aggregator at a time. A MAC client is served by a single aggregator at a time.

- The Link Aggregation Control function binds ports to aggregators within a system. The control function aggregates links, binds the system ports to an appropriate aggregator, and monitors conditions to determine if a change in aggregation is needed. Network managers can manually provide link aggregation control by manipulating the link aggregation state variables (for example, keys). You can also use LACP to automatically determine, configure, bind, and monitor link aggregation.

- LACP uses peer exchanges across links to continually determine the aggregation capability of the links and provide the maximum level of aggregation capability between a pair of systems.

- Frame ordering is maintained for certain sequences of frame exchanges between MAC Clients. The distributor ensures that all frames of a conversation pass to a single port. The collector passes frames to the MAC client in the order they are received from the port. The collector can select frames received from the aggregated ports. Because the

frames are not ordered on a single link, this guarantees that frame ordering is maintained for all conversations.

- Conversations move among ports within an aggregation for load balancing and for maintaining availability if a link fails.

- The standard does not impose a particular distribution algorithm on the distributor. The Virtual Services Platform 9000 uses the MLT traffic distribution algorithm. For more information on the Virtual Services Platform 9000 traffic distribution algorithm, see MLT traffic distribution algorithm on page 10.

- Each port is assigned a unique, globally administered MAC address.

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

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

Split MultiLink Trunking

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

SMLT overview

Split MultiLink Trunking is an option that improves Layer 2 (bridged) resiliency by providing for the addition of Virtual Services Platform 9000 failure redundancy with subsecond failover, on top of all standard MLT link failure protection and flexible bandwidth scaling functionality. Use Split MultiLink Trunking to connect a device which supports some form of link aggregation, be it a switch or a server, to two distinct separate SMLT endpoints or switches. These SMLT devices form a Switch Cluster and are referred to as an interswitch trunk (IST) Core Switch pair.

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

You can form SMLT connections through single links from the Switch Cluster to the edge connection, Single Link SMLT (SLT), standard MLTs, or MLTs with LACP. Optionally, SMLT links can have VLACP enabled as well. You can mix these various link connections. Within the same Switch Cluster, you can configure both SMLT and RSMLT to allow a mixture of both Layer 2 and Layer 3 VLANs. For more information about various SMLT configurations, see SMLT topologies on page 19.

Split MultiLink Trunking networks do not need to use RSTP or MSTP to enable loop-free triangle topologies because SMLT inherently avoids loops due to its superior enhanced link aggregation protocol. The loop-free link is accomplished by having two aggregation switches appear as a single device to edge switches, which dual-home to the aggregation switches. The aggregation switches interconnect using an interswitch trunk, exchanging addressing and state information (permitting rapid fault detection and forwarding path modification). Split MultiLink Trunking is designed for Layer 2 network connectivity, but you can configure in Layer 3 networks by working with VRRP or RSMLT Layer 2 edge.

For more information about SMLT configuration, see *Technical Configuration Guide for Switch Clustering using Split-MultiLink Trunking (SMLT), NN48500-518* and *SMLT and RSMLT Technical Configuration Guide*. For more information about these technical configuration guides, go to the Avaya Web site http://support.avaya.com.

SMLT advantages

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

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

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

SMLT, MSTP, and RSTP

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

- no load sharing exists over redundant links
- network convergence is slow in case of failure

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

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

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

SMLT topologies

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

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

SMLT and interswitch trunking

The following illustration shows an SMLT configuration with a pair of Virtual Services Platform 9000 systems as aggregation systems (E and F) and four separate Virtual Services Platform 9000 systems as user access switches (A, B, C, and D).
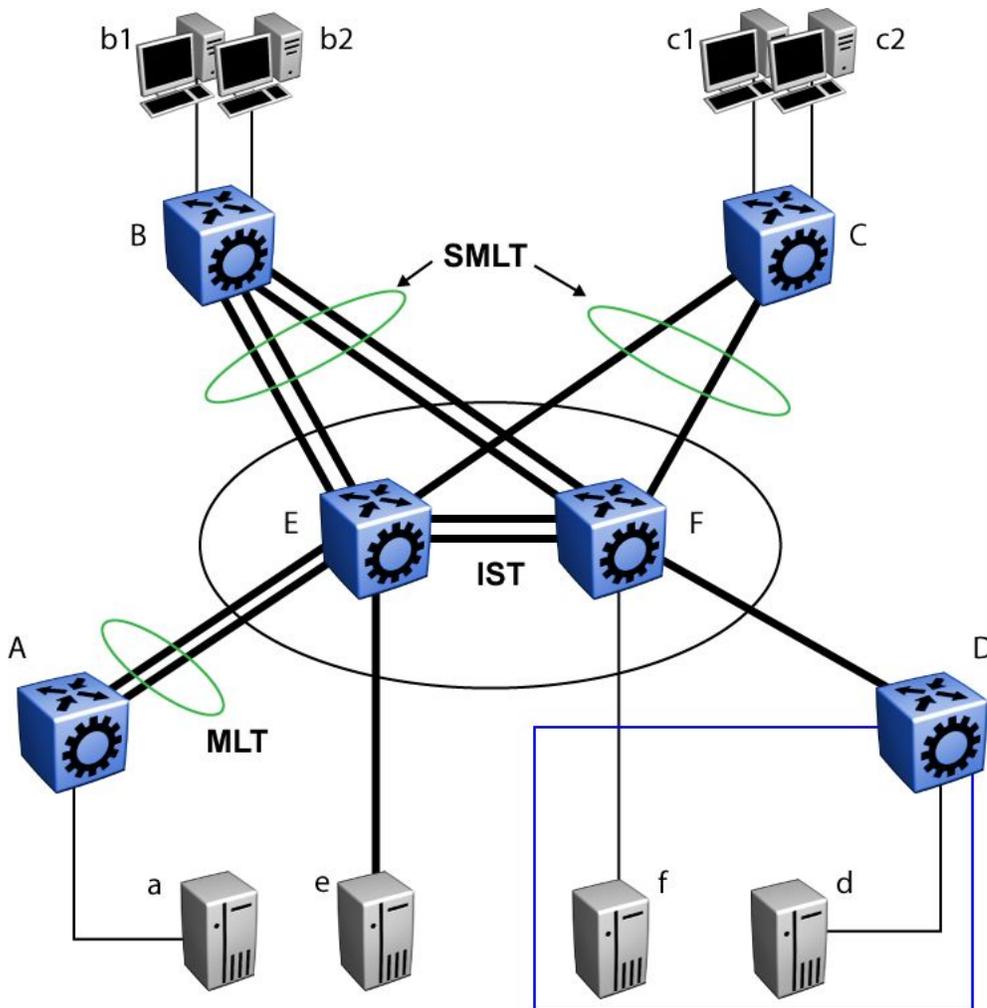
Figure 2: SMLT aggregation switches and operations

You must connect SMLT aggregation switches through an interswitch trunk. For example, user access switches B and C connect to the aggregation systems through multilink trunks split between the two aggregation systems. As shown in Figure 2: SMLT aggregation switches and operations on page 20, the implementation of SMLT only requires two SMLT-capable aggregation systems. You must connect these Virtual Services Platform 9000 systems through an interswitch trunk.

Aggregation switches use the interswitch trunk to do the following:

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

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

If the SMLT interfaces from one of the IST peers are down, all MAC entries and IP Address Resolution Protocol (ARP) entries learned from SMLT interfaces are unchanged. This behavior

is different than on Avaya Ethernet Routing Switch 8600, where all corresponding MAC and ARP entries are converted to be learned from the IST interface.

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

🛈 Important:
Avaya recommends that you do not configure CP-Limit on the IST MLT.

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

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

Other SMLT aggregation switch connections

The example shown in Figure 2: SMLT aggregation switches and operations on page 20 includes end stations that connect to each of the Virtual Services Platform 9000 systems. In this example, a, b1, b2, c1, c2, and d are clients and printers, while e and f are servers and routers.

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

SMLT environment traffic flow rules

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

- If a packet is received from an interswitch trunk port, it is not forwarded to an active SMLT group in order to prevent network loops.

- After a packet is received, the system performs a look-up on the forwarding database. If an entry exists, and if the entry was learned locally from the SMLT or through the IST as a remote SMLT, it is forwarded to the local port (the packet must not be sent to the IST for forwarding unless there is no local connection). Unknown and Broadcast packets flood out all ports that are members of this VLAN.

- For load sharing purposes in an SMLT configuration, the Virtual Services Platform 9000 obeys the MLT traffic distribution algorithm.

SMLT traffic flow examples

The following traffic flow examples are based on Figure 2: SMLT aggregation switches and operations on page 20.

Example 1: Traffic flow from a to b1 or b2

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

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

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

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

Example 3: Traffic flow from a to d

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

Example 4: Traffic flow from f to c1/c2

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

SMLT and IST traffic flow example

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

The following illustration shows the network topology.
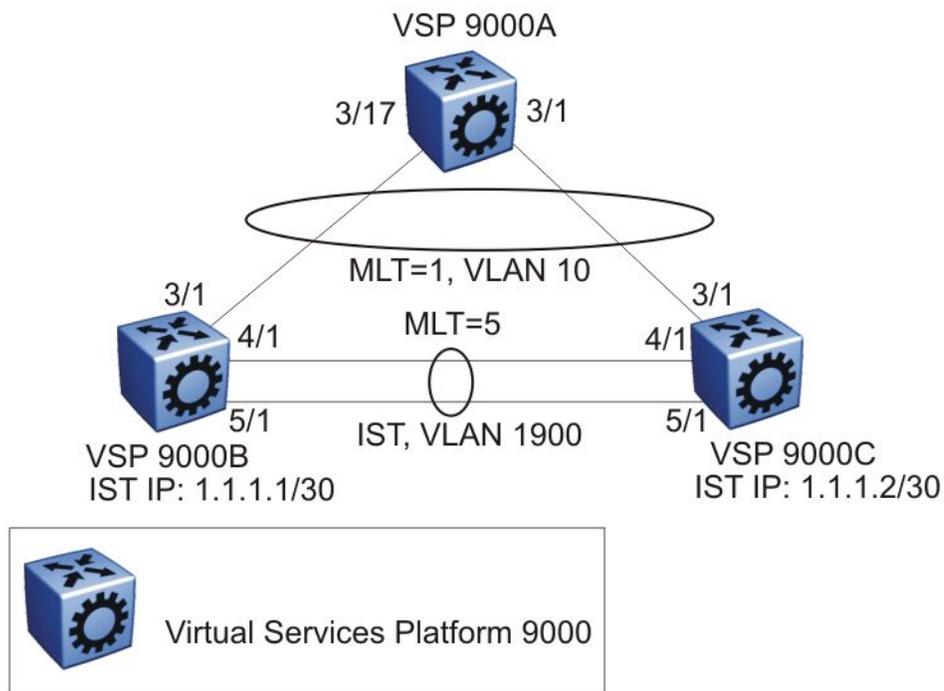
Figure 3: Network topology for traffic flow example

After a packet arrives at node C destined for 00:E0:7B:B3:04:00, if the SMLT Remote status is true, the Virtual Services Platform 9000 tries to send the packet to MLT-1 rather than through the interswitch trunk. Traffic rarely traverses the interswitch trunk unless there is a failure. If this same packet arrives at B, it is then forwarded to MLT-1 on the local ports.

SMLT-on-Single-CPU feature

The SMLT-on-Single-CPU feature reroutes traffic around SMLT aggregation systems with a single-failed CP module to avoid packet loss. This improves SMLT failover behaviors for single CP module configurations. You can also implement SMLT-on-Single-CPU on dual CP chassis independent of SMLT design. Because single CP module configurations do not benefit from standard CP module redundancy, in the rare event that a CP module failure occurs on the aggregation system, this enhancement forces the interface modules offline and activates network redundancy configurations quickly.

SMLT and LACP support

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

With LACP the Virtual Services Platform 9000 provides a standardized external link aggregation interface to third-party vendor IEEE 802.3ad implementations. This protocol extension provides dynamic link aggregation mechanisms. Only dual-home devices benefit from this enhancement.

Advantages of this protocol extension include the following:

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

Supported SMLT/LACP scenarios

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

Unsupported SMLT/LACP scenarios

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

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

SMLT and IP routing

This section describes SMLT and IP routing interactions.

-
-

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.

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

SMLT and VRRP BackupMaster

The VRRP BackupMaster acts as an IP router for packets destined for the logical VRRP IP address. The system directly routes all traffic to the destined subnetwork and not through Layer 2 switches to the VRRP master. This avoids a potential limitation in the available IST bandwidth.

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

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

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

Follow these guidelines if you use VRRP BackupMaster with SMLT:

- The VRRP virtual IP address and the VLAN IP address cannot be the same.

- Configure the hold-down timer for VRRP to a value that is approximately one hundred and fifty percent of the Interior Gateway Protocol (IGP) convergence time to allow the IGP enough time to reconverge following a failure. For example, if OSPF takes 40 seconds to reconverge, configure the hold-down timer to 60 seconds.

- Stagger the hold-down timers with ARP requests so the Virtual Services Platform 9000 does not run ARP and SMLT at the same time causing excess CPU load. For example, if one node has the hold-down timer configured for 60 seconds, you can configure the other hold-down timer to 65 seconds.

- Enable hold-down times on both VRRP sides (Master and BackupMaster).

SMLT and SLPP

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

Do not enable SLPP Rx on IST MLT ports. You can enable SLPP on the non-IST ports on an IST core switch pair. The Virtual Services Platform 9000 does not support SLPP in an SMLT that uses LACP.

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

For more information about SLPP fundamentals and configuration, see *Avaya Virtual Services Platform 9000 Administration, NN46250-600*.

Virtual Link Aggregation Control Protocol

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

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

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

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

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

For example, in , 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.
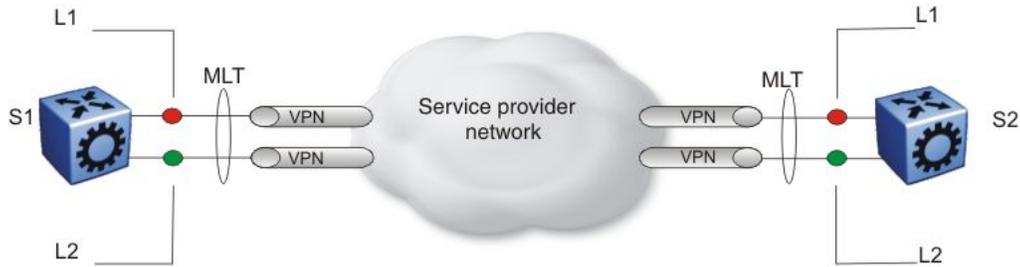
Figure 4: Problem description (1 of 2)

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.
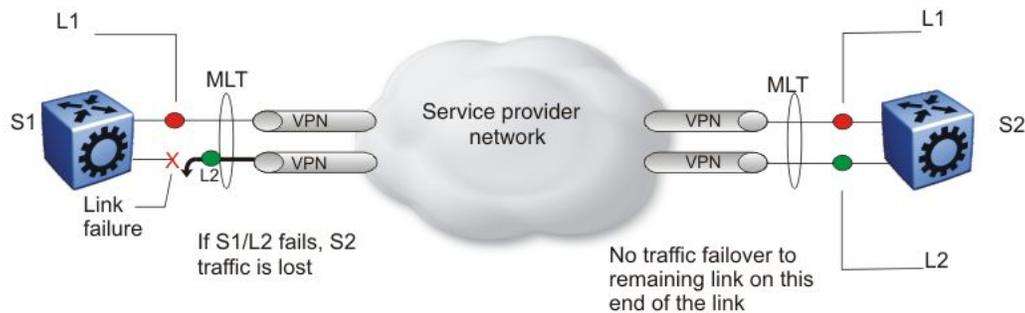
Figure 5: Problem description (2 of 2)

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

If used in conjunction with SMLT, VLACP switches traffic around entire network devices before Layer 3 protocols detect a network failure, thus minimizing network outages.

VLACP timers and sub-100-ms core convergence

Virtual Services Platform 9000 software uses the following VLACP timers:

  • fast periodic timer—100 to 20 000 ms; default 200 ms

  • slow periodic timer—10 000 to 30 000 ms; default 30 000 ms

Virtual Services Platform 9000 software can attain sub-100-ms failover time. The VLACP fast periodic timer is reduced to a minimum value of 30 ms to enable sub-100-ms convergence (period of timeout x fast periodic timer = 3 x 30 ms = 90 ms). Sub-100-ms convergence guarantees ultra-fast convergence for critical business and multimedia applications.

Though you can configure Virtual Services Platform 9000 ports with a VLACP fast periodic timer of less than 30 ms, this configuration is not stable in scaled networks, and therefore not supported. Avaya recommends using VLACP with a fast periodic timer of 30 ms and a timeout of 3 as the fastest timers to achieve sub-100-ms failover. In addition, Avaya only supports the

same minimum fast periodic timer configuration of 3 times 30 ms after you enable High Availability mode, no matter how scaled the environment.

VLACP considerations

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

- Avaya recommends that you use VLACP with long timers on MLT IST links.

- If a VLACP-enabled port does not receive a VLACP Data Unit (VLACPDU), it must enter the disabled state. There are occasions where a VLACP-enabled port does not receive a VLACPDU but remains in the forwarding state. To avoid this situation, ensure that the VLACP configuration at the port level is consistent. You must either enable or disable both sides of the point-to-point connection.

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

Link aggregation configuration considerations

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

MLT with LACP configuration considerations

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

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

1. Assign a numeric key to the ports you want to include in the LAG.

2. Configure the LAG for aggregation.

3. Enable LACP on the port.

4. Create an MLT and assign to it the same key as in step 1.

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

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

After a multilink trunk is configured with LACP, you cannot add or delete ports or VLANs manually without first disabling LACP.

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:

- After you set the LACP system ID for SMLT, configure the same LACP SMLT system ID on both aggregation switches to avoid loss of data. Avaya recommends that you configure the LACP SMLT system ID to be the base MAC address of one of the aggregate switches, and that you include the SMLT-ID. 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, Avaya recommends that you set the multicast packets-per-second value to 6000 pps.

- Avaya recommends that you do not enable LACP on ISTs to avoid unnecessary processing. Use VLACP if an optical network between the SMLT core switches requires a failure detection mechanism.

- Enabling multicast MLT distribution on a square-mesh SMLT can cause loss of streams going through the diagonal links. This results in traffic distributed improperly. Avaya recommends that you disable the diagonal links.

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

with its SMLT-ID. The administrator must ensure that the same value for the system ID is configured on both of the SMLT Core Aggregation Switches.

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

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

The LACP System ID is the base MAC address of the switch, which is carried in Link Aggregation Control Protocol Data Units (LACPDU). If two links interconnect two switches that run LACP, each switch knows that both links connect to the same remote device because the LACPDUs originate from the same System ID. If you enable the links for aggregation using the same key, LACP can dynamically aggregate them into an MLT LAG.

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

However, if the switch that owns that Base MAC address restarts, the IST goes down, and the other switch reverts to using its own Base MAC address as the LACP System ID. This action causes all edge switches that run LACP to think that their links are connected to a different switch. The edge switches stop forwarding traffic on their remaining uplinks until the aggregation can reform (which can take several seconds). Additionally, after the restarted switch comes back on line, the same actions occur, thus disrupting traffic twice.

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

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.

In contrast to MLT, MLT with LACP does not expect Bridge Protocol Data Unit (BPDU) to be replicated over all ports in the trunk group. Therefore, you must enter the `ntstg disable` command to disable the parameter on the Spanning Tree Group (STG) for LACP-based link aggregation. For more information about this command, see *Avaya Virtual Services Platform 9000 Configuration — VLANs and Spanning Tree, NN46250-500*.

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

> **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

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

- LACP priority on page 31
- LACP keys on page 31
- LACP timers on page 31
- LACP modes on page 32

LACP priority

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

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

LACP keys

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

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

LACP timers

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

Changes to LACP timer values at the global level are reflected on all ports. However, you can change the LACP timer values for each port level. After you change an LACP timer globally, this value is set on all ports. The global timer value overwrites the local port value irrespective of the LACP state. You must configure port values that differ from the global values.

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

- fast periodic timer—200 to 20 000 milliseconds (ms); default 1000 ms
- slow periodic timer—10 000 to 30 000 ms; default 30000 ms
- aggregation-wait timer—200 to 2000; default 2000

You cannot aggregate a link if it does not receive an LACPDU for a period of timeout x slow periodic time = 3 x 30 seconds = 90 seconds. If you use the fast periodic time, the timeout

period is 3 x 1000 ms = 3 seconds. You must make timer changes to all ports participating in link aggregation and to the ports on the partnering node.

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

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

LACP modes

LACP uses two active and passive modes.

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

Link aggregation scaling

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

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

Feature	Number
Aggregation groups • 802.3ad aggregation groups (LAG) • multilink trunk groups	512
Maximum number of links (ports) for each MLT group	16

Chapter 4: LACP configuration using EDM

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

> **Important:**
> Virtual Services Platform 9000 does not support standby mode for aggregation groups larger than eight ports.

Configuring global LACP parameters

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

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

Table 2: Variable definitions

Variable	Value
Enable	Enables or disables LACP globally.
SystemPriority	Configures the system priority for all LACP enabled aggregators and ports. The default value is 32 768.
FastPeriodicTime	Configures the number of milliseconds between periodic transmissions that use short timeouts. Sets this value to all LACP-enabled ports. The range is 200–20 000. 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 10 000–30 000. The default value is 30 000.

Variable	Value
SlowPeriodicTimeOper	Displays the operating value of the slow periodic timer on the port. The default value is 30 000.
AggrWaitTime	Configures the number of milliseconds to delay aggregation to allow multiple links to aggregate simultaneously. The range is 200–2000. The default value is 2000.
AggrWaitTimeOper	Displays the operating value of the aggregate wait timer on the port. The default value is 2000.
TimeoutScale	Configures the value used to calculate timeout time from the periodic time. All LACP-enabled ports get the same value from this setting. The range is 2–10. The default value is 3.
TimeoutScaleOper	Displays the operating value of the timeout scale on the port. The default value is 3.
SmltSysId	Configures the LACP system ID for SMLT. Use this optional parameter only for SMLT situations. You must configure the same LACP SMLT system ID on both aggregation switches to avoid loss of data. Avaya recommends that you configure the SmltSysId so it matches the base MAC address of one of the chassis. The default value is 00:00:00:00:00:00.

1. In the navigation tree, open the following folders: **Configuration, VLAN**.

2. Double-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**.

Configuring LACP parameters

Prerequisites

🛈 **Important:**

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

Configure LACP parameters to manage LACP information.

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

Table 3: Variable definitions

Variable	Value
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).

Variable	Value
ActorAdminKey	The current administrative value of the key for the aggregator. The administrative key value can differ from the operational key value, which is a 16–bit read-write value. This key needs to match the LAG key. The default value is 0.
ActorOperKey	The current operational value of the key for the aggregator. The administrative key value can differ from the operational key value, which is a 16–bit read-only value. The meaning of particular key values is of local significance.
PartnerSystemID	The six octet read-only MAC address value consisting of the unique identifier for the current protocol partner of this aggregator. A value of zero indicates that there is no known partner. If the aggregation is manually configured, the value is assigned by the local system.
PartnerSystemPriority	The two octet read-only value that indicates the priority value associated with the partner system ID. If the aggregation is manually configured, this system priority value is a value assigned by the local system.
PartnerOperKey	The current operational value of the key for the aggregator current protocol partner, which is a 16-bit read-only value. If the aggregation is manually configured, the value is assigned by the local system.

1. In the navigation tree, open the following folders: **Configuration, VLAN**.

2. Double-click **MLT/LACP**.

3. Click the **LACP** tab.

4. Double-click on a field to edit it.
   You cannot edit some fields.

5. Click **Apply**.

Configuring LACP on a port

Configure LACP on a port to enable LACP.

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

Table 4: Variable definitions

Variable	Value
AdminEnable	Enables LACP status for the port. The default value is false.
OperEnable	Displays the operational status of LACP for the port. The default value is false.
FastPeriodicTime	Specifies the number of milliseconds between periodic transmissions using short timeouts for all LACP enabled ports. The range is 200–20000. The default value is 1000.
FastPeriodicTimeOper	Displays the operating value of the fast periodic timer on the port. The default value is 1000.
SlowPeriodicTime	Specifies the number of milliseconds between periodic transmissions using long timeouts for all LACP enabled ports. The range is 10000–30000. The default value is 30000.
SlowPeriodicTimeOper	Displays the operating value of the slow periodic timer on the port. The default value is 30000.
AggrWaitTime	Specifies the number of milliseconds to delay aggregation to allow multiple links to aggregate simultaneously. The range is 200–2000. The default value is 2000.
AggrWaitTimeOper	Displays the operating value of the aggregate wait timer on the port. The default value is 2000.
TimeoutScale	Assigns the value used to calculate timeout time from the periodic time. Set this value to all LACP enabled ports. The range is 2–10. The default value is 3.
TimeoutScaleOper	Displays the operating value of the timeout scale on the port. The default value is 3.
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.

Variable	Value
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 distinction is made between the values of these parameters for an aggregator and the ports that are associated with it; that is, the protocol is described in terms of the operation of aggregation within a single system. However, the managed objects provided for the aggregator and the port both allow management of these parameters. The result of this is to permit a single piece of equipment to be configured by management to contain more than one system from the point of view of the operation of link aggregation. This can be of particular use in the configuration of equipment that has limited aggregation capability.
ActorAdminKey	Specifies the current administrative value of the key for the aggregator. The administrative key value can differ from the operational key value, which is a 16-bit read-write value in the range of 0–65535. This key needs to match the LAG key.
ActorOperKey	Displays the current operational value of the key for the aggregator. The administrative key value can differ from the operational key value, which is a 16-bit read-only value. The meaning of particular key values is of local significance.
SelectedAggID	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.
AttachedAggID	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.

Variable	Value
ActorPort	Displays the port number locally assigned to the aggregation port. The port number is communicated in LACPDUs as the Actor_Port. This value is read-only.
ActorPortPriority	Specifies the priority value assigned to this aggregation port. This 16-bit value is read-write in the range of 0–65535. The default value is 32768.
ActorAdminState	Specifies a string of eight bits, corresponding to the administrative values as transmitted by the actor in LACPDUs, by selecting check boxes. The values are • 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 allow administrative control over the values of LACP_Activity, LACP_Timeout, and aggregation. This attribute value is read-write.
ActorOperState	Displays a string of eight bits, corresponding to the current operational values of Actor_State as transmitted by the actor in LACPDUs. This attribute value is read-only. The default value is lacpActive.
PartnerAdminSystemPriority	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,

Variable	Value
	and PartnerAdminPortPriority, to achieve manually configured aggregation. The default value is 0.
PartnerOperSystemPriority	Displays a two octet read-only value indicating the operational value of priority associated with the partner system ID. The value of this attribute can contain the manually configured value carried in PartnerAdminSystemPriority if there is no protocol partner. The default value is 0.
PartnerAdminSystemID	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. The default value is 00:00:00:00:00:00.
PartnerOperSystemID	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. The default value is 00:00:00:00:00:00.
PartnerAdminKey	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. The default value is 0.
PartnerOperKey	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. The default value is 0.
PartnerAdminPort	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–

Variable	Value
	65535. The assigned value is used, along with the value of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminKey, and PartnerAdminPortPriority, to achieve manually configured aggregation. The default value is 0.
PartnerOperPort	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. The default value is 0.
PartnerAdminPortPriority	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. The default value is 0.
PartnerOperPortPriority	Displays the priority value assigned to this aggregation port by the partner. The value of this attribute can contain the manually configured value carried in PartnerAdminPortPriority if there is no protocol partner. This 16 bit value is read-only. 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.

1. In the Device Physical View tab, select a port.

2. Open the following folders: **Configuration, Edit, Port**.

3. Double-click **General**.

4. Click the **LACP** tab.

5. To enable LACP on the port, select the **AdminEnable** check box.

6. Configure the remaining parameters as required.

7. Click **Apply**.

Chapter 5: LACP configuration using ACLI

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

MultiLink Trunking (MLT) with 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.

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).

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

Important:
After you globally set 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. 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.

- Standby mode for LACP aggregation groups of larger than eight ports is not supported.
- Configuring global LACP parameters on page 44
- Configuring LACP on a port on page 45
- Viewing LACP configuration information on page 47

Configuring global LACP parameters

Prerequisites

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

🛈 Important:

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

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

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

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

Table 5: Variable definitions

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

Configure LACP globally:

```
lacp
```

Configuring LACP on a port

Prerequisites

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

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

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

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

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

Table 6: Variable definitions

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

Variable	Value
`partner-state <0-255| 0-255>`	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 • Act = true • Agg = true specify a value of 0x05 (bitmap = 00000101).
`partner-system-id <0x00:0x00:0x00:0x00:0 x00:0x00>`	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.
`timeout-scale <2-10>`	Configures a timeout scale for this port. The default value is 3. To set this option to the default value, use the **default** operator with the command. The LACP timeout is equal to the slow periodic time or fast periodic time multiplied by the timeout-scale, depending how you set the timeout-time variable.
`timeout-time {long| short}`	Configures the timeout to either long or short. To set this option to the default value, use the **default** operator with the command.

Configure LACP on a port:

```
lacp
```

Viewing LACP configuration information

Prerequisites

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

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

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

Table 7: Variable definitions

Variable	Value
`actor-admin interface gigabitethernet [vid <1-4084> [{slot/port[-slot/port] [,...]}]`	Shows LACP actor administrative information for all interfaces or the specified interface. • **<1-4084>** is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • **{slot/port[-slot/port] [,...]}** is the port or port list.
`actor-oper interface gigabitethernet [vid <1-4084> [{slot/port[-slot/port] [,...]}]`	Shows LACP actor operational information for all interfaces or the specified interface. • **<1-4084>** is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • **{slot/port[-slot/port] [,...]}** is the port or port list.
`extension interface gigabitethernet [vid <1-4084> [{slot/port[-slot/port] [,...]}]`	Shows LACP timer information for all interfaces or the specified interface. • **<1-4084>** is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • **{slot/port[-slot/port] [,...]}** is the port or port list.

Variable	Value
`interface [gigabitethernet [vid <1-4084> [{slot/port[- slot/port][,...]}]] [mlt <64-6399>}]`	Shows all LACP port configuration information for all interfaces or the interface you specify. • **<1-4084>** is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • **{slot/port[-slot/port] [,...]}** is the port or port list. • **<64-6399** is the interface index of the mlt.
`partner-admin interface [gigabitethernet] [vid <1-4084> [{slot/port[-slot/ port][,...]}]`	Shows LACP partner administrative information for all interfaces or the specified interface. • **<1-4084>** is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • **{slot/port[-slot/port] [,...]}** is the port or port list.
`partner-oper interface [gigabitethernet] [vid <1-4084> [{slot/port[-slot/ port][,...]}]`	Shows LACP partner operational information for all interfaces or the specified interface. • **<1-4084>** is the VLAN ID or list of VLAN IDs to show only ports attached to a particular VLAN. The range is 1–4084. Can be written as a range of VLAN IDs. • **{slot/port[-slot/port] [,...]}** is the port or port list.

View LACP configuration information:

```
show lacp
```

Chapter 6: 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, interswitch, or LACP trunk

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

ⓘ Important:

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

Use the data in the following table to help you configure the MultiLink/LACP tab

Table 8: Variable definitions

Variable	Value
Id	Specifies a unique value for this MLT.
PortType	Specifies the value to access or trunk port. If the aggregatable field is set to enable, this field is read-only. The default value is access.
Name	Configures the name given to the MLT.
PortMembers	Assigns ports to the MLT. 10Base-T, 100Base-TX, 100Base-FX, and Gigabit Ethernet ports support MLT. All ports in an MLT must have the same settings for speed and duplex, but can have different media types. All untagged ports must belong to the same STG. For Virtual Services Platform 9000 modules, up to eight same-type ports can belong to a single MLT. If the aggregatable field is set to enable, this field becomes read-only.

Variable	Value
	🛈 **Important:** Ensure that all ports that belong to the same MLT/LACP group use the same port speed, for example, 1 Gbit/s, even if autonegotiation is enabled.
VlanIds	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 • normalMLT (default) • istMLT • splitMLT
RunningType	Specifies the MLT running type.
IfIndex	Specifies the interface of the trunk
MulticastDistribution	Enables or disables multicast distribution for this trunk.
ClearLinkAggregate	Clears the link aggregate, disabling and re-enabling the trunk.
DesignatedPort	Specifies the designated port of this trunk.
Aggregatable	Enables or disables link aggregation. The default value is disable.
AggOperState	Specifies the aggregation state of the trunk.
AggTimeOfLastOperChange	Specifies the time value since the interface entered its current operational state.
PeerPortMembers	Specifies the peer ports connected to the local ports of this trunk.

1. In the navigation tree, open the following folders: **Configuration, VLAN**.

2. Double-click **MLT/LACP**.

3. Click the **MultiLink/LACP Trunks** tab.

4. Click **Insert**.

5. In the **Id** box, type the ID number of the MLT.

6. In the **PortType** section, select **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 **VlanIds** box, click the (**...**) button,

12. In the **VlanIds** dialog box, select the desired VLANs.

13. Click **Ok**.

14. In the **MltType** section, select the MLT type.
    If you select splitMLT, the Smltid box becomes editable.

15. In the **MulticastDistribution** box, select **enable** or **disable**.

16. In the **Aggregatable** box, select **enable** or **disable**.

17. Click **Insert**.
    The MLT is added to the MultiLink/LACP Trunks tab in the MLT_LACP box.

Adding ports to an MLT

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

1. In the navigation tree, open the following folders: **Configuration, VLAN**.

2. Double-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.

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

Table 9: Variable definitions

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

1. In the navigation tree, open the following folders: **Configuration, VLAN**.

2. Double-click **SMLT**.

3. Click the **SMLT Info** tab.

Configuring an IST MLT

Configure an IST MLT to add an IST to MLT_LACP.

1. In the navigation tree, open the following folders: **Configuration, VLAN**.

2. Double-click **MLT/LACP**.

3. Click the **MultiLink/LACP Trunks** tab.

4. Click **Insert**.

5. In the **PortMembers** box for the IST MLT, click the **(...)** button.

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

7. Click **Ok**.

8. In the **MltType** box, select **istMLT**.

9. In the **PortType** box, select **trunk**.

10. Configure the remaining parameters as required.

11. Click **Insert**.

12. Type the peer IP address and the VLAN ID in the box.

13. Click **enable**.

14. Click **Apply**.

15. In the Device Physical View tab, select all IST MLT ports.

16. In the navigation tree, open the following folders: **Configuration, Edit, Port**

17. Double-click **General**.

18. Click the **CP Limit** tab.

19. In the **CpLimitEnable** option for each port, select **disable**.

20. Click **Apply**.

Editing an IST

Edit an IST to make changes to an IST MLT.

Use the data in the following table to help you configure the MultiLink/LACP Trunks-IstMlt tab.

Table 10: Variable definitions

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

1. In the navigation tree, open the following folders: **Configuration, VLAN**.

2. Double-click **MLT/LACP**.

3. Click the **MultiLink/LACP Trunks** tab.

4. Select the IST.

5. Click **IstMlt**.

6. In the **PeerIp** box, type the peer IP address.

7. In the **VlanId** box, enter a VLAN ID.

8. In the **SessionEnable** box, select either **enable** or **disable**.

9. Click **Apply**.

Chapter 7: MLT and SMLT link aggregation configuration using the ACLI

This chapter describes how to configure and manage link aggregation using the ACLI, including MultiLinkTrunking (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 an MLT

Prerequisites

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

Perform this procedure to configure or create an MLT on the Avaya Virtual Services Platform 9000.

Use the data in the following table to use the `mlt <1-512>` command.

Table 11: Variable definitions

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

Configure an MLT:

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

Adding ports to an MLT LAG

Prerequisites

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

Perform this procedure to add ports to an MLT LAG.

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

Table 12: Variable definitions

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

Adding ports to an MLT LAG:

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

Removing ports from an MLT LAG

Prerequisites

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

🛈 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.

Remove ports from an MLT LAG.

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

Table 13: Variable definitions

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

1. Disable the ports:
   ```
   shutdown port {slot/port[-slot/port][,...]}
   ```
2. Remove ports from an MLT LAG:
   ```
   no mlt <1-512> member {slot/port[-slot/port][,...]}
   ```

Configuring LACP on an MLT

Prerequisites

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

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

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

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

Table 14: Variable definitions

Variable	Value
enable	Enables LACP on the MLT interface.
key <0-512>	Sets the LACP aggregator key for a specific MLT. *0-256* is the LACP actor admin key.
system-priority *<0-65535>*	Sets the LACP system priority for a specific MLT. *0-65535* is the system priority.

Configure LACP on an MLT:

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

Creating an SMLT from an existing MLT

Prerequisites

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

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

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

Create an SMLT from an existing MLT:

```
smlt
```

Creating an interswitch trunk

Prerequisites

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

Create an interswitch trunk from an existing MLT.

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

Table 15: Variable definitions

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

Create an interswitch trunk:

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

Enabling an interswitch trunk

Prerequisites

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

Perform this procedure to enable an interswitch trunk.

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

Table 16: Variable definitions

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

Enable an interswitch trunk:

```
ist enable
```

Deleting an interswitch trunk

Prerequisites

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

🛈 **Important:**

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

Delete an interswitch trunk.

1. Disable an interswitch trunk:

   ```
   no ist enable
   ```

2. Delete an interswitch trunk:

   ```
   no ist peer-ip
   ```

Example of deleting an interswitch trunk

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

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

2. Disable the IST:

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

3. Delete the IST:

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

Viewing the MLT port calculated by the MLT hash algorithm

Prerequisites

- You must log on to the Global Configuration mode in ACLI.
- View the MLT port calculated by the MLT hash algorithm to view the MLT port calculated by the new distribution algorithm.

View the MLT port calculated by the MLT hash algorithm to obtain information about the MLT port calculated by the new distribution algorithm.

Use the data in the following table to use the `sys hash-calc` command.

Table 17: Variable definitions

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

View the MLT port calculated by the MLT hash algorithm:

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

The source and destination addresses cannot have the same value.

Viewing information about collision errors

Prerequisites

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

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

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

Table 18: Variable definitions

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

View information about collision errors:

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

Viewing information about Ethernet errors

Prerequisites

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

ⓘ Important:
The IMAC columns refer to internal MAC address 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.

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

Table 19: Variable definitions

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

View information about Ethernet errors:

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

Viewing all ports configured for SMLT

Prerequisites

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

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

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

Table 20: Variable definitions

Variable	Value
mlt	Displays SMLT information for the MLT interface.

View all ports configured for Single Port SMLT:

```
show smlt mlt
```

Chapter 8: 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 on page 65
- Configuring VLACP on a port on page 65

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.

Use the data in this table to use the VLACP Global tab.

Variable	Value
VlacpEnable	Enables VLACP globally.

1. In the navigation tree, open the following folders: **Configuration, VLAN**.
2. Double-click **MLT/LACP**.
3. Click the **VLACP Global** tab.
4. Select the **VlacpEnable** check box.
5. Click **Apply**.

Configuring VLACP on a port

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

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

Table 21: Variable definitions

Variable	Value
AdminEnable	Enables VLACP for the port.
OperEnable	Displays the operational status of VLACP for the port. The default value is false.
FastPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using short timeouts. Set all LACP enabled ports the same value from this setting. The range is 10–20000. The default value is 200.
SlowPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using long timeouts. Set all LACP enabled ports the same value from this setting. The range is 10000–30000. The default value is 30000.
Timeout	Specifies the timeout control value. Specify long or short timeout. The default value is long.
TimeoutScale	Assigns the value used to calculate timeout time from the periodic time for all VLACP enabled ports. Timeout = PeriodicTime x TimeoutScale. The range is 2–10. The default value is 3.
EtherType	Specifies the VLACP protocol identification. The ID is in hexadecimal format. The default value is 8103.
EtherMacAddress	Specifies the multicast MAC address exclusively used for VLACPDUs.
PortState	Displays the VLACP port state.

1. Select a port.

2. In the navigation tree, open the following folders: **Configuration, Edit, Port**.

3. Double-click **General**.

4. In the Port General tab, click the **VLACP** tab.

5. Select the **AdminEnable** check box.

6. Configure the remaining parameters as required.

7. Click **Apply**.

Chapter 9: VLACP configuration using ACLI

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

Configuring VLACP on a port

Prerequisites

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

🛈 Important:

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

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.

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

Table 22: Variable definitions

Variable	Value
`enable`	Enables VLACP for this port.
`ethertype` *0x1-0xffff*	Configures the VLACP protocol identification for this port. Enter the type in hex format.
`fast-periodic-time` *<100-20000>*	Configures the fast periodic time (in milliseconds) for this port.
`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.

Variable	Value
slow-periodic-time *<10000-30000>*	Configures the slow periodic time (in milliseconds) for a specific port type.
timeout {long|short}	Configures the port to use the long or short timeout: • **long** sets the port to use the timeout-scale value multiplied by the slow periodic time. • **short** sets the port to use the timeout-scale value multiplied by the fast periodic time. For example, if you specify a short timeout, set the timeout-scale value to 3, and the fast periodic time to 400 ms, the timer expires within 1000 to 1200 ms. To set this option to the default value, use the **default** operator with the command.
timeout-scale <2-10>	Configures a timeout scale for this port used to calculate the timeout. The default value is 3. To set this option to the default value, use the **default** operator with the command.

Configure VLACP on a port:

```
vlacp
```

Viewing the VLACP port configuration

Prerequisites

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

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

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

Table 23: Variable definitions

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

View the VLACP port configuration for all interfaces:

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

Enabling or disabling VLACP globally

Prerequisites

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

🛈 **Important:**

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

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

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

Table 24: Variable definitions

Variable	Value
`enable`	Enables VLACP globally. To disable VLACP globally, use the no form of this command: `no vlacp enable`.

Enable VLACP globally:

```
vlacp enable
```

OR

Disable VLACP globally:

```
no vlacp enable
```

Chapter 10: MLT network topology and configuration reference

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

- • Example 1: Switch-to-switch MLT on page 71
- • Example 2: Switch-to-server MLT on page 72
- • Example 3: Client/Server MLT on page 72

Example 1: Switch-to-switch MLT

The following illustration shows two multilink trunks (MLT1 and MLT2) connecting VSP 9000 C to switches VSP 9000 A and VSP 9000 B.
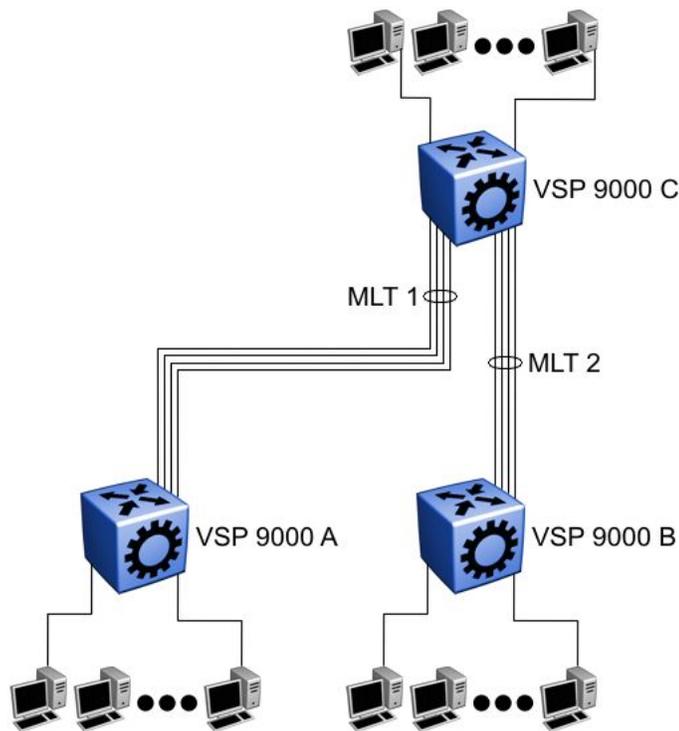
Figure 6: Switch-to-switch multilink trunks configuration

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

bandwidth limitations, you can create a multilink trunk to supply the additional bandwidth required to improve performance, as well as providing physical link layer redundancy.

Example 2: Switch-to-server MLT

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

The following illustration shows a typical switch-to-server multilink trunk configuration.
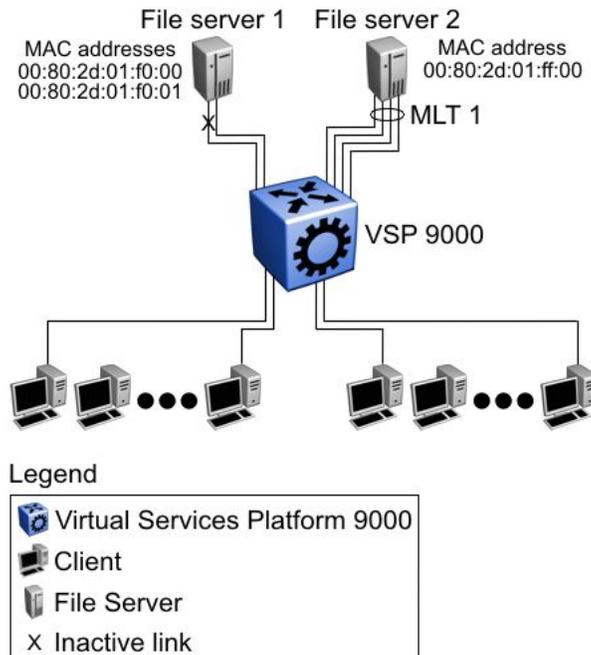
Figure 7: Switch-to-server multilink trunk configuration

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

Example 3: Client/Server MLT

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

The following illustration shows how you can use multilink trunks in a client/server configuration.
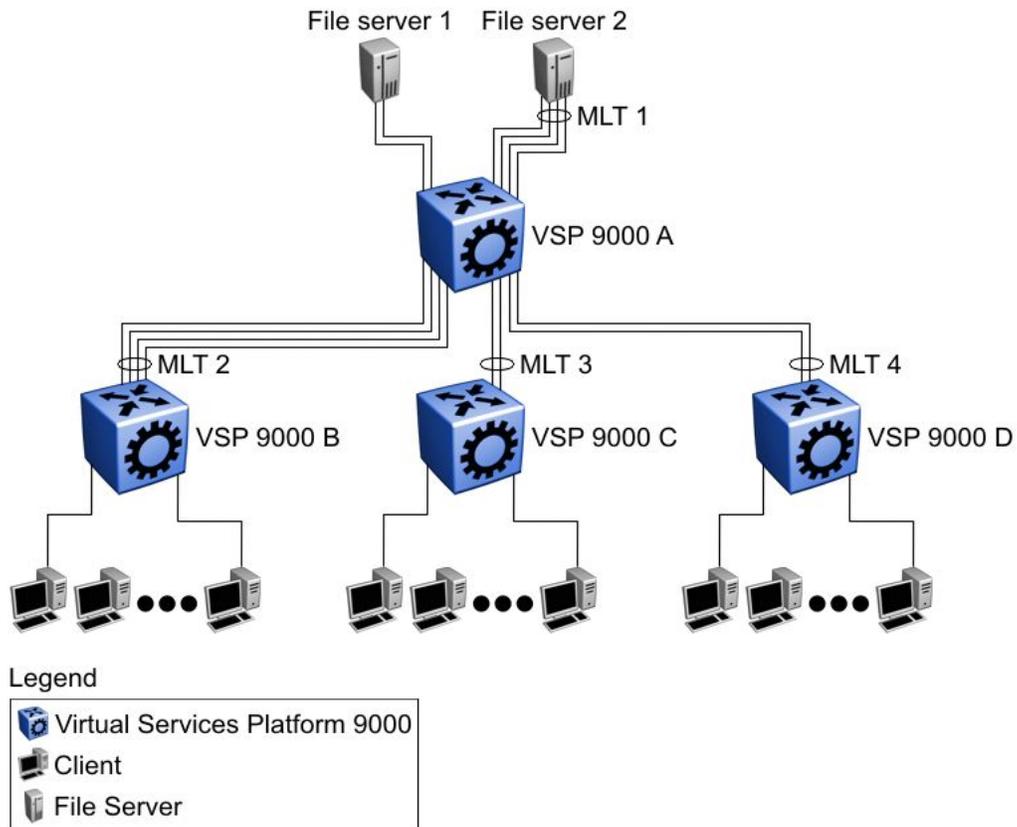
Figure 8: Client/server multilink trunk configuration

Chapter 11: Configuration examples using ACLI

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

MultiLink Trunking configuration example

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

The following illustration shows you an MLT within a VLAN used to carry user traffic.
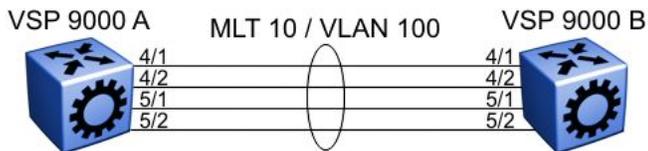
Figure 9: MLT within a VLAN

Configuring VSP 9000 A

Prerequisites

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

Perform this procedure to configure MLT 10 on VLAN 100 on VSP 9000 A.

1. Create VLAN 100:

   vlan create 100 type port-mstprstp 1

2. Create multilink trunk 10:

```
mlt 10 enable
```

3. Assign ports to MLT 10:

```
mlt 10 member 4/1, 4/2, 5/1, 5/2
```

4. Assign VLAN 100 to MLT 10:

```
mlt 10 vlan 100
```

Configuring VSP 9000 B

Prerequisites

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

Perform this procedure to configure MLT 10 on VLAN 100 on VSP 9000 B.

1. Create VLAN 100:

```
vlan create 100 type port-mstprstp 1
```

2. Create multilink trunk 10:

```
mlt 10 enable
```

3. Assign ports to MLT 10:

```
mlt 10 member 4/1, 4/2, 5/1, 5/2
```

4. Assign VLAN 100 to MLT 10:

```
mlt 10 vlan 100
```

Next steps

MultiLink Trunking with Link Aggregation Control Protocol configuration example

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.
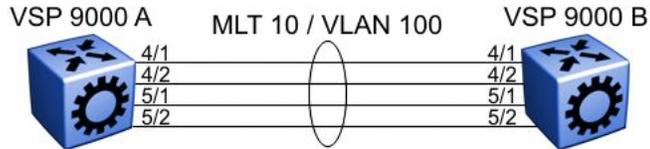
Figure 10: MLT within a VLAN

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

Configuring VSP 9000 A

Prerequisites

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

Configure VSP 9000 A with MLT with LACP to connect to VSP 9000 B.

1. Create VLAN 100:

   vlan create 100 type port-mstprstp 1

2. Add ports to VLAN 100:

   vlan members add 100 4/1,4/2,5/1,5/2

3. Enter Interface Configuration mode for ports 4/1, 4/2, 5/1, 5/2:

   interface gigabitethernet 4/1,4/2,5/1,5/2

4. Configure the LACP key:

   lacp key 10

5. Enable link aggregation:

   lacp aggregation enable

6. Enable LACP:

   lacp enable

7. Return to Global Configuration mode:

   exit

8. Create and enable MLT 10:

   mlt 10 enable

9. Enter Interface MLT Configuration mode:

   interface mlt 10

10. Associate LACP key to MLT:

```
lacp key 10
```

11. Activate LACP on MLT 10:

```
lacp enable
```

Configuring VSP 9000 B

Prerequisites

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

Configure VSP 9000 B with MLT with LACP to connect to VSP 9000 A.

1. Create VLAN 100:

```
vlan create 100 type port-mstprstp 1
```

2. Add ports to VLAN 100:

```
vlan members add 100 4/1,4/2,5/1,5/2
```

3. Enter Interface Configuration mode:

```
interface gigabitethernet 4/1,4/2,5/1,5/2
```

4. Configure the LACP key:

```
lacp key 10
```

5. Enable link aggregation:

```
lacp aggregation enable
```

6. Enable LACP:

```
lacp enable
```

7. Return to Global Configuration mode:

```
exit
```

8. Create and enable MLT 10:

```
mlt 10 enable
```

9. Enter Interface MLT Configuration mode:

```
interface mlt 10
```

10. Associate LACP key to MLT:

```
lacp key 10
```

11. Activate LACP on MLT:

```
lacp enable
```

———

Chapter 12: Customer service

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

Navigation

- Getting technical documentation on page 81
- Getting product training on page 81
- Getting help from a distributor or reseller on page 81
- Getting technical support from the Avaya Web site on page 82

Getting technical documentation

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

Getting product training

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

Getting help from a distributor or reseller

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

Getting technical support from the Avaya Web site

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

Index

T

V