Data Center Server Access Solution Guide

NN48500-577
Date: July 2010
Version: 1.1

© 2010 Avaya Inc.
All Rights Reserved.

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

Documentation disclaimer
Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of this documentation unless such modifications, additions, or deletions were performed by Avaya. End User agree to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

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

Warranty
Avaya provides a limited warranty on this product. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this product, while under warranty, is available to Avaya customers and other parties through the Avaya Support Web site: http://www.avaya.com/support
Please note that if you acquired the product from an authorized reseller, the warranty is provided to you by said reseller and not by Avaya.

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

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

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

Trademarks
The trademarks, logos and service marks ("Marks") displayed in this site, the documentation(s) and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the documentation(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party. Avaya is a registered trademark of Avaya Inc. All non-Avaya trademarks are the property of their respective owners.

Downloading documents
For the most current versions of documentation, see the Avaya Support. Web site: http://www.avaya.com/support.

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

Abstract

The Data Center Server Access Solution Guide provides the basic information necessary to design and configure Server connectivity in the Data Center. Detailed descriptions of NIC teaming along with Avaya Ethernet infrastructure design parameters are covered. The solution guide provides high level descriptions of the technologies, recommended best practices, and configuration examples to make implementation easy.

Revision Control

No	Date	Version	Revised by	Remarks
1	03/16/2009	1.0	D. DeBacker	Initial version of the solution guide

Table of Contents

Figures

Tables

Document Updates

July 30, 2010

1. Introduction

The drive towards virtualization and consolidation in the Data Center causes many changes in the network design. This Solution Guide addresses the needs of the server access layer and the technologies that make virtual server deployments more reliable, simple, scalable, and efficient.

Computing virtualization allows customers to run multiple applications on a single physical server. This computing virtualization creates new challenges of ensuring that the physical server doesn't become a single point of failure in terms of hardware, network connectivity, or storage access. If one of the network interfaces fails, an alternate path must exist to maintain network connectivity. This same requirement applies to the storage connectivity path albeit not as critical as the production network connectivity.

In the case of an entire physical server failure, the virtualization management tools must detect and move the Virtual Machine to another physical server. This virtualization paradigm changes the way applications are deployed and moved across physical servers within and between Data Centers. A unique demand on the network infrastructure in terms of reliability and performance now play a key role in the overall network infrastructure requirements.

The Data Center Server Access layer must be scalable and able to support multiple applications/workloads running on a single physical server as well as provide the redundancy required for Enterprise customers expecting 7 x 24 x 365 availability.

In a typical server, there are anywhere from two to eight or more Network Interface Cards (NICs) providing customers with different options of connectivity depending on the bandwidth needs of the server and reliability requirements. If there is a need to increase the network bandwidth, multiple NICs are combined together through Link Aggregation thus providing increased connectivity. At the same time, highly virtualized servers need the redundancy as the loss of a single network interface will affect numerous users across multiple departments or business units. The impact of a loss of connectivity on a virtualized server is greater than network connectivity loss for a non-virtualized server. As a result, redundancy takes on a new level of importance.

The following highlights areas that significantly improve the Server Access Layer:

- NIC Teaming
- Link Aggregation
- Horizontal Stacking
- Switch Clustering

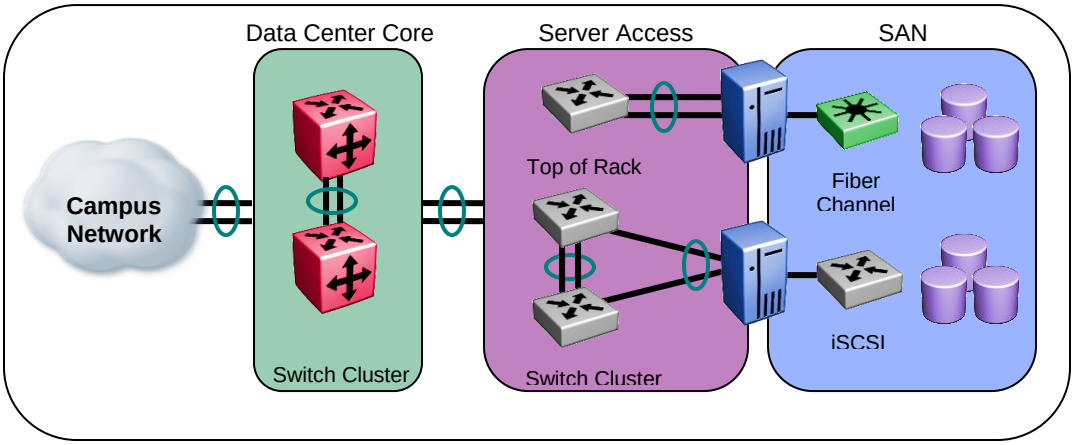

Figure 1.1: Data Center Architecture Example

2. NIC Teaming Fundamentals

NIC teaming is the process of grouping together (binding) several physical NICs into one single logical NIC. Teaming requires software that is typically provided by the NIC vendor such as Intel, Broadcom, HP, or others. In the case of VMware, the NIC teaming feature is part of the ESX platform and no additional software is required while in the Microsoft Hyper-V environment, NIC teaming is not part of the Hyper-V platform, but instead is done at the physical server level.

Let us briefly look at how NIC teaming works and the different modes supported by the NIC teaming software.

The server shown in Figure 2.1 has 2 NICs and no teaming software but still provides a level of dual homing. In this scenario, each NIC has its own IP address. A user can access the server through NIC1 via 1.1.1.1. If this adapter fails, the connection between the server and user is lost. The user can connect to the server again through the other NIC via 1.1.1.2, but this is not a seamless process and interrupts the user's server access experience. This solution does not provide an adequate level of redundancy.

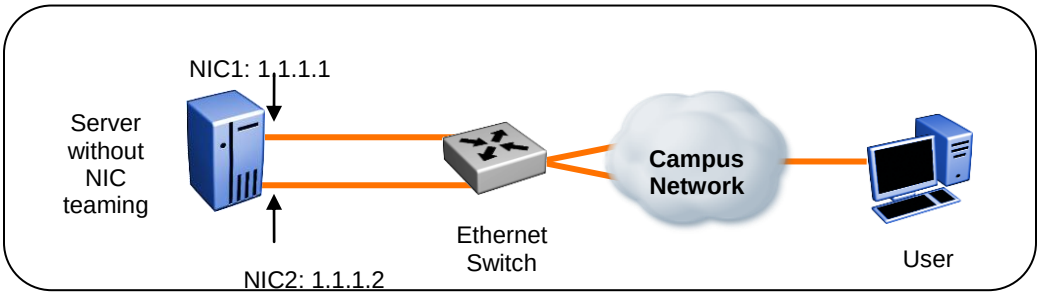

Figure 2.1: Server without NIC Teaming

In contrast, with NIC teaming software, a single virtual NIC adapter is created with a single IP address. As shown in Figure 2.2, the user accesses the server via one IP address. In the case of a failure of one of the NICs, the user's connectivity to the server will not be impacted as traffic will automatically and seamlessly fail over to the remaining connection. This solution now provides both the resiliency and performance required for Enterprise Data Centers.

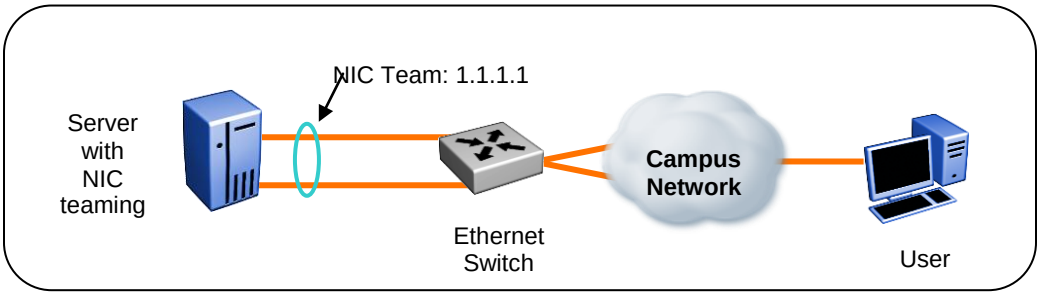

Figure 2.2: Server with NIC Teaming

3. Ethernet Infrastructure Fundamentals

The Ethernet infrastructure for the Server Access Layer within the Data Center is critical to providing the required resiliency, performance, and manageability to meet Enterprise requirements. For the purposes of this Solution Guide, we will look at a few main concepts that provide this infrastructure and also highlight the unique features that differentiate Avaya.

- Link Aggregation
  - Multilink Trunking
  - 802.3ad
- Horizontal Stacking
- Switch Clustering
  - SMLT
  - SLT

3.1 Link Aggregation

In order to increase both resiliency and bandwidth from the Server to the Ethernet Switch, Avaya recommends implementing link aggregation. Avaya supports two options for link aggregation on the Ethernet switches, MultiLink Trunking (MLT) and 802.3ad.

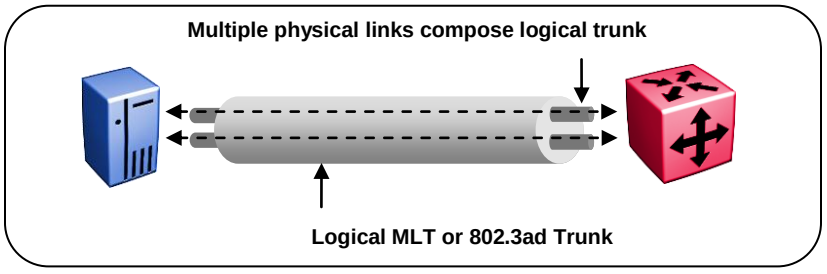

Figure 3.1: Link Aggregation

MultiLink Trunking (MLT) provides the ability to group multiple physical links into a single logical link. MLT automatically increases bandwidth to/from the Server by utilizing all the physical links in a logical group. A failure of any physical links results in automatic sub-second failover of the traffic to the remaining links within the MLT group. After the failed link has been repaired, recovery of that link back into the MLT group is also accomplished in sub-second time.

Distributed MultiLink Trunking (DMLT) adds the flexibility to terminate the physical links of the trunk group on different switches within a stack or on different I/O modules within a chassis. This feature significantly increases the resiliency of the uplinks. A failure of the switch or I/O module on which one of these physical links terminates will not cut off communication from the Server to the Server Access Layer.

The IEEE standard for link aggregation is 802.3ad. This standard uses a Link Aggregation Control Protocol (LACP) to aggregate multiple physical links into a single logical link aggregation group (LAG) from the Ethernet switch. Adherence to the IEEE standard will help to ensure interoperability of link aggregation between different vendor equipment (switches, NICs, etc.).

Although IEEE 802.3ad-based link aggregation and MLT provide similar services, MLT is statically defined, whereas IEEE 802.3ad-based link aggregation is dynamic and provides added functionality through the Link Aggregation Control Protocol (LACP). LACP dynamically detects whether links can be aggregated into a link aggregation group and does so when links become available. IEEE 802.3ad was designed for point-to-point link aggregation only. The Avaya Ethernet switches support the formation of

802.3ad link aggregation groups across different switches in a stack or different I/O modules in a chassis to provide additional resiliency in the event on a switch failure in a stack or an I/O module failure in a chassis.

Switch Model	Links per Group	Groups per Switch or Stack	802.3ad Support	LACP-MLT Scaling	LACP-SMLT Scaling
ERS 8600 Legacy Modules	8	32	Rls.3.7	32	32
ERS 8600 R / RS Modules	8	128	Rls 5.0	128	128
ERS 8300	4	31*	Rls 4.1	31	31
ERS 5000	8	32	Rls 4.1	32	Future

* Up to seven Fast Ethernet groups and/or 31 Gigabit groups

Table 3.1: Link Aggregation Scaling

3.2 Horizontal Stacking

The ERS 5000 switches support a resilient stacking architecture, using Avaya's FastStack technology with a shortest path algorithm used for stacking, allows for the most efficient use of bandwidth across the stack. A failure in any unit of the stack will not adversely affect the operation of the remaining units in the stack. Replacement of the failed switch is easy with the Auto-Unit Replacement feature which allows for a new unit to be put into the stack and automatically get the right software image and configuration. The entire process can be done live without any resets of the entire stack necessary.

The unique advantage Horizontal Stacking brings to the Server Access Layer is the ability to leverage the high bandwidth stack capabilities of the ERS 5000 series switches. The ERS 5500 switches support 80Gbps per switch stack bandwidth, for a maximum bandwidth of 640Gbps in a full stack of eight switches. The ERS 5600 series switches support 144Gbps per switch of stack bandwidth, for a maximum bandwidth of 1.1Tbps in a full stack of eight switches.

Another big advantage of Horizontal Stacking is the reduction in the amount of uplinks from the Server Access Layer to the Data Center Core. With individual Top of Rack switches, each rack would normally have one or two uplinks to the core. With a Horizontal Stack, the number of uplinks is reduced significantly, while still providing the flexibility to add uplink bandwidth capacity as required.

The Horizontal Stack can be created by using the various lengths of stacking cables available from Avaya. These cables come in lengths of 1, 1.5, 3, 10, and 16.4 feet to provide the flexibility needed when connecting switches between cabinets in the Data Center. Once the Horizontal Stack has been cabled up correctly, it is recommended to renumber the units in the stack. The base unit should be the leftmost unit in the stack and identified as unit #1. Moving from left to right, number the units in the stack as #2, #3, and #4. If a stack of greater than four units is created, continue with the numbering in sequential order.

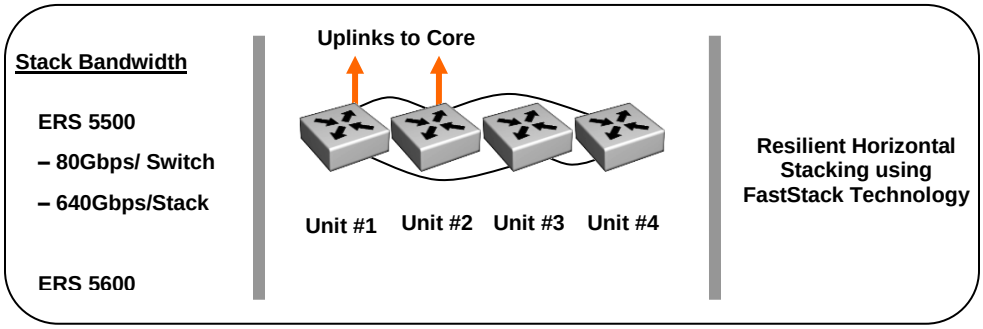

Figure 3.2: ERS 5000 Horizontal Stacking

3.3 Switch Clustering

Switch Clustering technology allows for dual homing of multiple links from the Server in an N-1 redundancy technique – all links active and passing traffic simultaneously. In the event of a link, switch, or module failure, Switch Clustering provides sub-second failover. The advent of Switch Clustering obsoletes the need for Spanning Tree protocol and its complexity.

With Switch Clustering, the aggregation switches appear as one logical device to the dual homed Server. All the intelligence of Switch Clustering rests in these aggregation switches and therefore, the technology is edge agnostic – meaning that any edge device that supports link aggregation can take advantage of Switch Clustering. The aggregation switches make use of an Inter Switch Trunk (IST) to exchange topology information, permitting rapid fault detection and forwarding path modification.

Switch Clustering also provides the ability to perform virtual hitless upgrades of the core switches (cluster). With all connections to the cluster dually attached, a single core can be taken out of service with minimal (sub-second) interruption to a portion of end user traffic. This switch/stack then can be upgraded and brought back into service. By performing the same function on the other switch/stack, after the upgraded switch/stack is back online, the entire cluster can be upgraded without taking a service outage and with minimal (sub-second) interruption to traffic flows on the network.

A vital feature of Switch Clustering is its ability to work transparently with any end device that supports some form of link aggregation. These end devices include 3rd party switches, servers, or appliances.

3.3.1 Switch Clustering Terminology

There are different design options to be considered with the deployment of Switch Clustering:

- Single Link Trunking (SLT)

  SLT is a port-based option allowing large-scale deployments of SLT from a single Switch Cluster. Every port (saving at least two for the IST) can be used for SLT groups terminating into the cluster, with each SLT group consisting of a maximum of two uplinks (one per core Ethernet Routing Switch). For most typical deployments, the ability to have two connections per server is more than sufficient bandwidth. The flexibility of the Avaya Ethernet switch solutions allows for uplinks ranging from 10 Mbps to 10 Gbps (uplinks within the same SLT group must be of the same media type and link speed).

- Split MultiLink Trunking (SMLT)

  The MLT-based SMLT option allows for increased scaling of the number of links within a single SMLT group. The number of links supported in an SMLT group is the same number of MLT links supported on the Ethernet Routing Switch platform being used for the Switch Cluster. The SMLT links can be spread across the Switch Cluster – usually in an even dispersion, but this is not an absolute requirement. One MLT group must be used to create the IST between the two switches used to form the Switch Cluster.

Both SMLT and SLT can be used simultaneously on all Switch Cluster configurations.

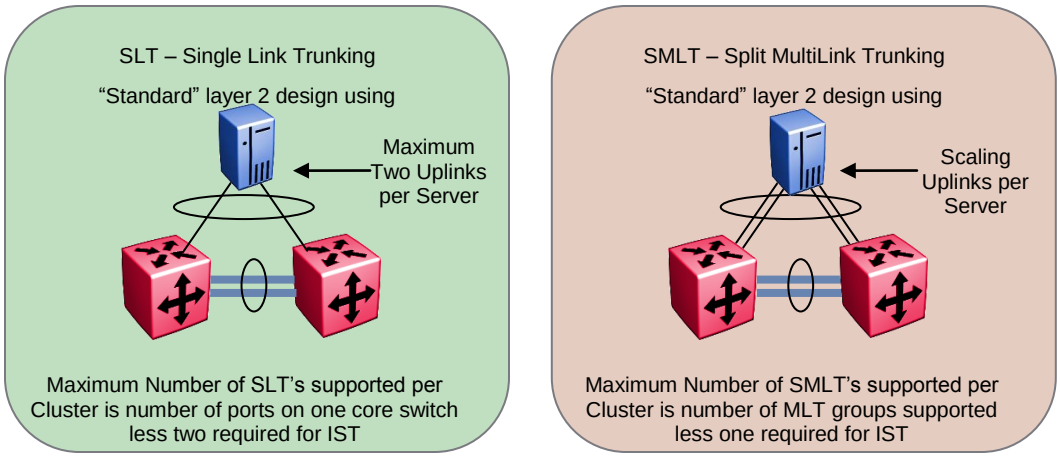

Figure 3.3: SLT and SMLT Terminology

Table 3.2 highlights the scaling capabilities of the various Ethernet Switch platforms with regard to MLT, SMLT, SLT capabilities.

Switch Model	Links per MLT Group	MLT Groups per Switch or Stack	MLT-based SMLT Groups			Port-based SLT Groups		
			Copper	Fiber (1GbE)	Fiber (10GbE)	Copper	Fiber (1GbE)	Fiber (10GbE)
ERS 8600 Legacy Modules	8	32	31	31	31	382	238	22
ERS 8600 R, RS Modules	8	128	127	127	127	382	238	22
ERS 8300	4	31	30	30	30	382	398	62
ERS 5000	8	32	31	31	31	382	94	14

Note: Advanced Software License required on ERS 8300 and ERS 5000 for SMLT

Table 3.2: MLT/SMLT/SLT Scaling Capabilities

3.3.2 Horizontal Stacking with Switch Clustering

The ERS 5000 series support Avaya Switch Clustering. By combining the resilient stacking architecture with the resiliency of Switch Clustering using Split MultiLink Trunking (SMLT) or Single Link Trunking (SLT), the overall Data Center design provides the highest level of resiliency for all dual connected devices. The ideal design includes installing a switch from each Horizontal Stack in a server rack. This would allow dual connections from the server to the Ethernet switches inside a rack – making cable installation, maintenance and troubleshooting much easier as all network connections are contained in the server rack.

An IST will be used to create the Horizontal Stack Switch Cluster. The IST will be comprised of an MLT of two to eight ports and can be Gigabit or 10Gigabit. Size and scale of the IST is dependent upon the number of servers and the speed the servers are connected to the Switch Cluster. Under normal operating conditions, the IST does not forward a significant amount of traffic as all connections are dual homed. The IST will carry traffic is any servers are single-homed or in the event of an uplink failure where all traffic must traverse a single switch/stack.

The IST will be distributed across multiple switches in the stack for added resiliency, with at least one IST member residing on the Base Unit of each stack. Please note that Spanning Tree must be manually disabled on the IST ports on the ERS 5000.

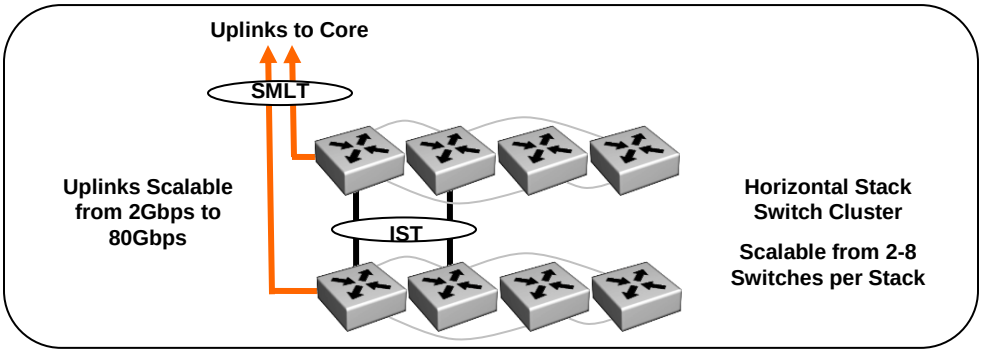

Figure 3.4: ERS 5500 Switch Clustering

3.4 General Recommendations

Autonegotiation

- Disable autonegotiation on both the server and the corresponding switch port and hard code the speed and duplex setting on each end. These ports are usually tightly controlled by the network administrator and therefore the capabilities of the device connected will not change. Disabling autonegotiation on these ports eliminates a variable in any troubleshooting activities for server connectivity.

Spanning Tree

- Disable Spanning Tree on the Ethernet Switch when connecting servers via MLT, DMLT, SMLT, or SLT. Spanning Tree is automatically disabled on ERS 8600 and ERS 8300 SMLT and SLT ports, but must be manually disabled on ERS 5000 SMLT and SLT ports. Spanning Tree must always be manually disabled on all ERS platforms when connecting servers via MLT or DMLT.
- Leave Spanning Tree enabled on all ports that are not in use to help prevent any unintentional loops being created by configuration mistakes or cables attached incorrectly. All non-uplink ports should be configured for Spanning Tree Fast Start.

VLANs and Routing

- In a two tier Data Center architecture (Core and Server Access), it is recommended to keep the Server Access as Layer 2 and perform routing at the Data Center Core. This provides the flexibility to create a Data Center utility model and extend the presence of VLANs wherever they are needed in the Data Center. This also allows a central point for all routing which reduces complexity and troubleshooting.
- For highly scaled Data Center environments, it may be necessary to extend Layer 3 to the Distribution (in a three tier model) or all the way to the Server Access layer. The flexibility of the Avaya infrastructure allows this without the need to totally re-architect the Data Center model.

Default Gateway Resiliency – VRRP

- VRRP provides redundancy for the server's default gateway and should be utilized for each VLAN configured that contain servers. Avaya has created an extension to VRRP that allows for local processing of traffic that would otherwise have to take an extra hop to get to the default gateway. VRRP Backup Master allows both core switches to forward and route traffic, creating an active-active environment for routing. This feature is extremely beneficial when implemented in conjunction with Switch Clustering.

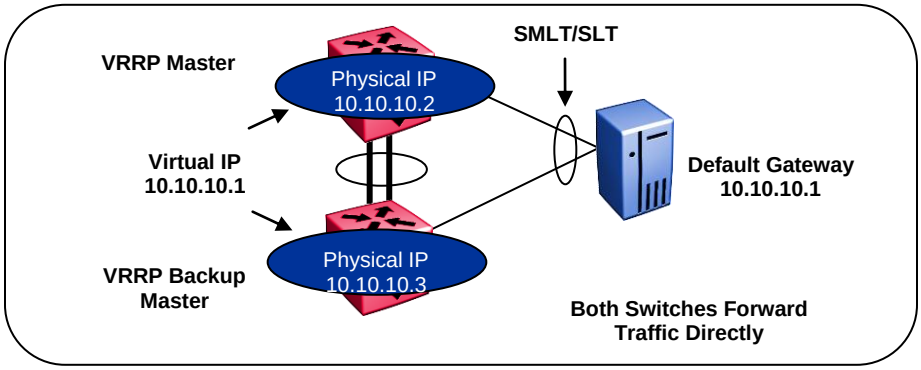

Figure 3.5: Default Gateway Resiliency – VRRP

- Enable VRRP and Backup Master on each VLAN
- For each IP subnet, stagger the VRRP Master between the two core Switch Cluster switches. By default, the VRRP priority is set to 100, therefore, configure the VRRP priority to 110 on the VRRP Master
- Leave VRRP priority at default (100) for VRRP Backup
- Do not configure the virtual address as a physical interface – use a third address
  - Physical IP address of VLAN on Switch 1 = x.x.x.2
  - Physical IP address of VLAN on Switch 2 = x.x.x.3
  - Virtual IP address of VLAN a = x.x.x.1
- Set Advertise Interval to 10 seconds
- Set Hold-Down Timer to 60 seconds
- Do not use the same VRID across multiple VLANs
- Do not enable VRRP Fast Advertise

Default Gateway Resiliency – RSMLT Layer 2 Edge

- RSMLT Layer 2 Edge is presently available on the ERS 8600 and ERS 8300 platforms. This feature provides redundancy for the server's default gateway and is a viable alternative to VRRP. Implementing RSMLT Layer 2 Edge instead of VRRP can provide several advantages; namely, RSMLT is only limited by the number of IP interfaces on the ERS 8600 or ERS 8300 where VRRP is limited to 250 instances; RSMLT requires significantly less control traffic and is much less intensive on CPU resources.

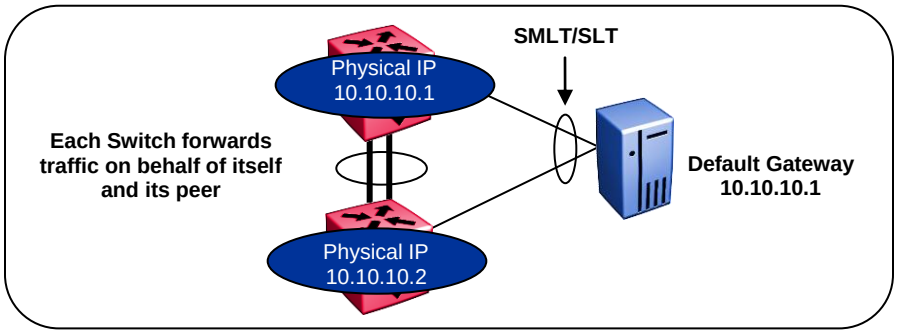

Figure 3.6: Default Gateway Resiliency – RSMLT Layer 2 Edge

- Please note that either VRRP or RSMLT Layer 2 Edge should be used, but not both simultaneously on the same VLAN
- Based on SMLT, so all SMLT rules apply
- Configured on a per VLAN basis
- VLAN must be routable and part of the SMLT links and IST link
- Hold up timer must be increased to 9999 (meaning infinity) so that the functioning switch is able to forward traffic indefinitely for a failed peer
- After enabling RSMLT Layer 2 Edge, make sure to save the configuration file on both peer switches

4. Broadcom NIC Teaming

Broadcom supports multiple different modes of NIC teaming. Each of these modes will be described briefly here, followed by the Avaya recommendations for which modes to use with various Avaya Ethernet switch solutions.

4.1 Smart Load Balancing

Smart Load Balancing (SLB) uses software to balance routable traffic among a team of two to eight adapters (the team must include at least one server adapter) connected to the same subnet. The software analyzes the transmit loading on each adapter and balances the rate across the adapters based on destination address. Adapter teams configured for SLB also provide the benefits of fault tolerance.

- SLB is switch independent and supports switch fault tolerance by allowing the teamed ports to be connected to more than one switch in the same LAN
- It takes advantage of all available bandwidth (No standby links)
- Only use Smart Load Balancing when connecting to switches without any link aggregation
- SLB does not load balance non-routed protocols such as NetBEUI and some IPX* traffic
- You can create an SLB team with mixed speed adapters. The load is balanced according to the lowest common denominator of adapter capabilities and the bandwidth of the channel
- On Windows systems, Receive Load Balancing is enabled by default

4.2 802.3ad using LACP

IEEE 802.3ad, which is also known as Link Aggregation Control Protocol (LACP), is an IEEE specification. The main advantage of LACP is it is a standard based implementation which helps in providing interoperability with switches from different vendors.

4.3 FEC/GEC Generic Trunking

The Generic Trunking (FEC/GEC)/802.3ad-Draft Static type of team is very similar to the Link Aggregation (802.3ad) type of team in that all adapters in the team need to be configured to receive packets for the same MAC address. The Generic Trunking (FEC/GEC)/802.3ad-Draft Static) type of team, however, does not provide LACP or marker protocol support. This type of team supports a variety of environments in which the adapter link partners are statically configured to support a proprietary trunking mechanism. For instance, this type of team could be used to support Lucent's OpenTrunk or Cisco's Fast EtherChannel (FEC) or Avaya's SMLT. Basically, this type of team is a light version of the Link Aggregation (802.3ad) type of team. This approach is much simpler, in that there is not a formalized link aggregation control protocol (LACP). As with the other types of teams, the creation of teams and the allocation of physical adapters to various teams are done statically through user configuration software. The Generic Trunking (FEC/GEC/802.3ad-Draft Static) type of team supports load balancing and failover for both outbound and inbound traffic which is provided by the logic that is implemented on the switches.

4.4 Avaya Recommendations

The following figures and tables highlight the different Ethernet Infrastructure topologies for the Server Access Layer along with the Avaya recommended NIC Teaming mode to be used.

Multilink Trunking / Link Aggregation to Modular / Standalone / Stack

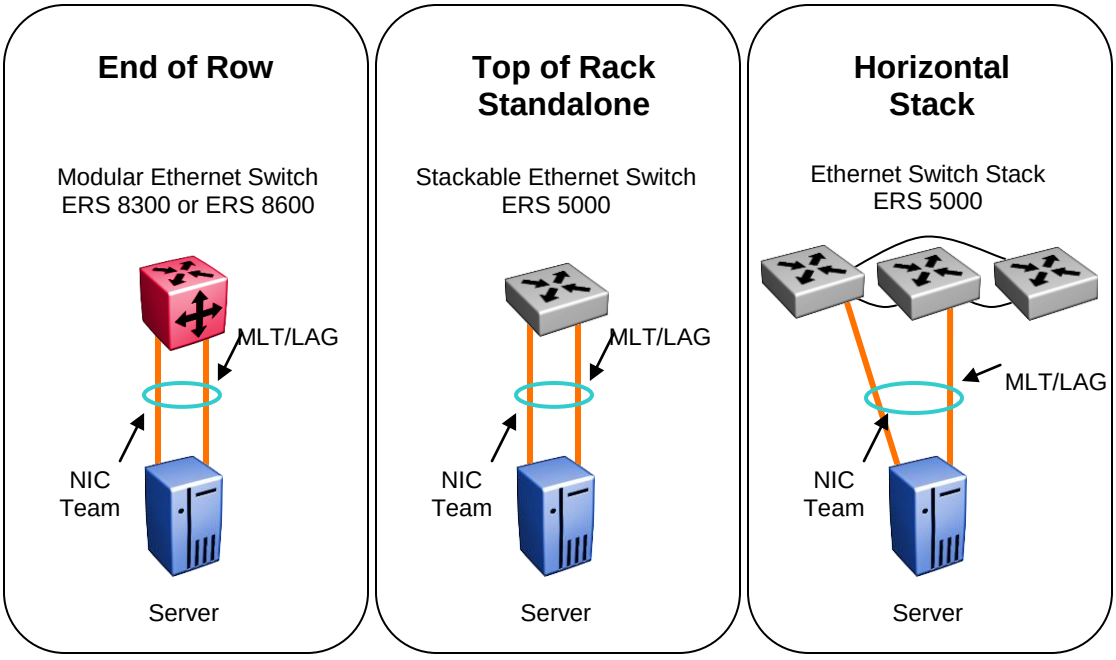

Figure 4.1: MLT/LAG with Broadcom NICs

Ethernet Infrastructure	ERS Platform	MLT/LAG Groups	Ports per MLT/LAG Group	NIC Teaming Mode
Standalone Switch/Stack using MLT to Server	ERS 5000	32	8	GEC Generic Trunking Or 802.3ad LACP
	ERS 8300	31	4	
	ERS 8600	128	8	

Notes:

- Total number of MLT/LAG groups equals total number of NIC teamed servers supported
- Do not use Smart Load Balancing with this Ethernet Infrastructure topology – it is not compatible with any type of Link Aggregation techniques.

Switch Clustering to Modular / Standalone / Stack

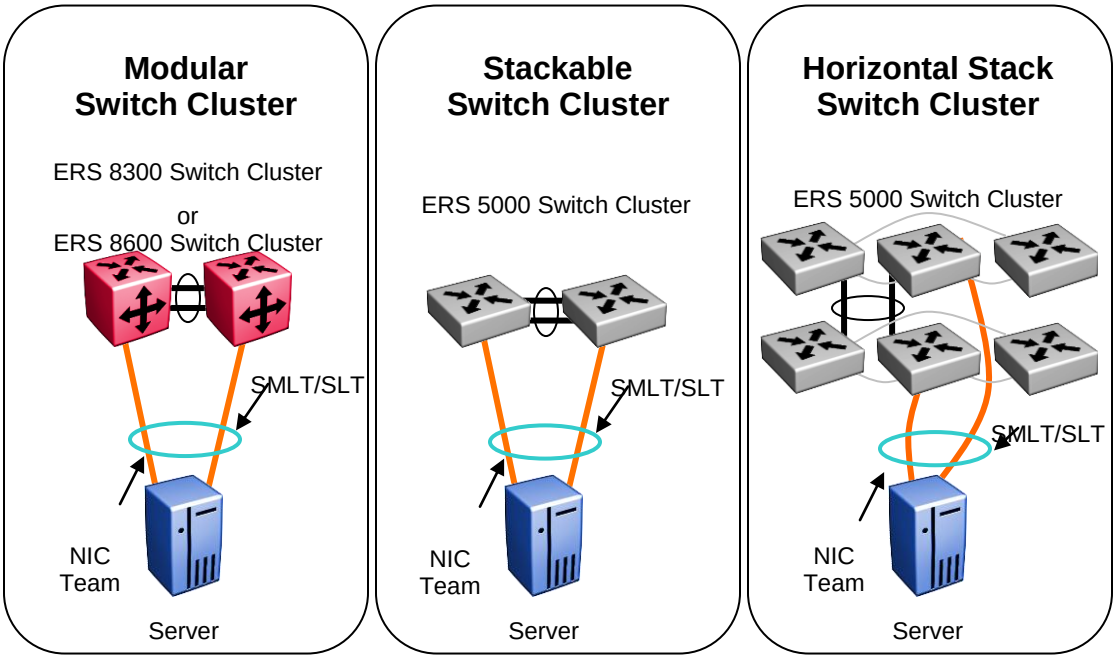

Figure 4.2: Switch Clustering with Broadcom NICs

Ethernet Infrastructure	ERS Platform	SMLT Groups	SLT Groups	NIC Teaming Mode
Switch Clustering to Server	ERS 5000	31	382	FEC/GEC Generic Trunking
	ERS 8300	30	382	FEC/GEC Generic Trunking Or 802.3ad LACP
	ERS 8600	127	382	

Note:

- SMLT and SLT can be used on the same Switch Cluster
- If using the ERS 8600, configure LACP SMLT-SYS-ID on each of the ERS 8600s
- Do not use 802.3ad Link Aggregation using LACP with ERS 5000 Switch Clusters as LACP over SMLT/SLT is not presently supported.
- Do not use Smart Load Balancing with this Ethernet Infrastructure topology – it is not compatible with any type of Link Aggregation techniques.

4.5 NIC Team Configuration Example

This section shows the details steps to configure the Broadcom NIC team for FEC/GEC Generic Trunking – which is the preferred NIC teaming mode to work with the Ethernet infrastructure topologies discussed in the previous section.

Broadcom Advanced Control Suite (BACS) is used to create the NIC team

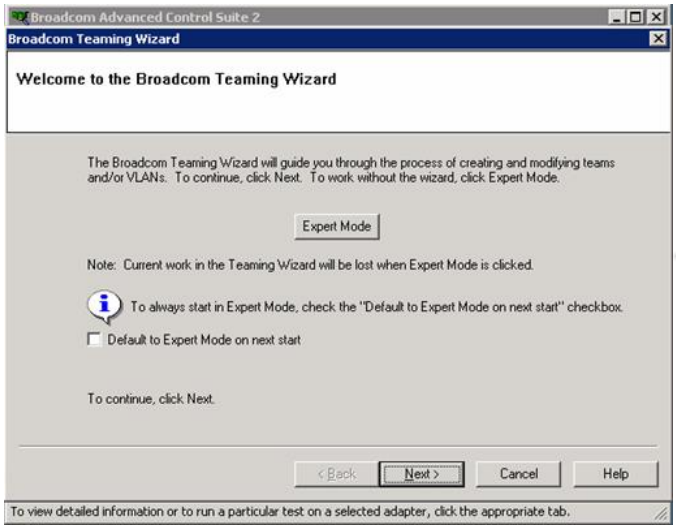

Broadcom Teaming wizard guides you through the process, step by step

NIC details

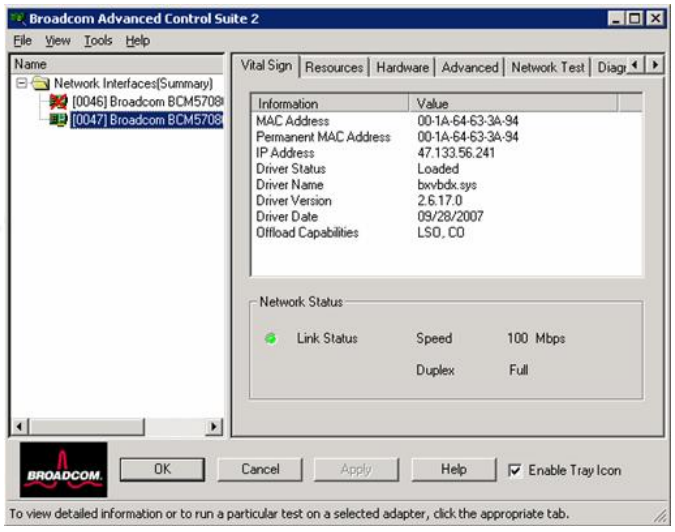

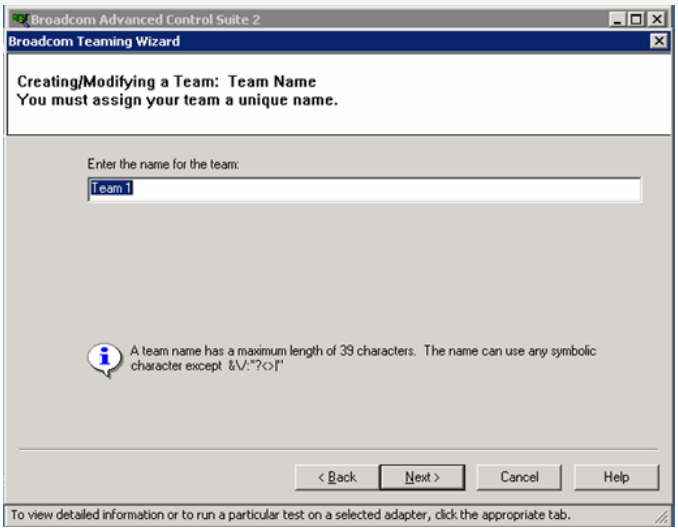

Create NIC Team name

Select FEC/GEC Generic Trunking

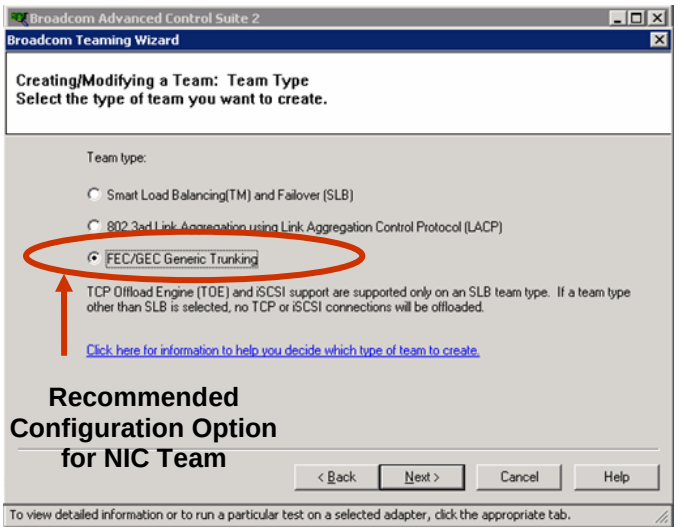

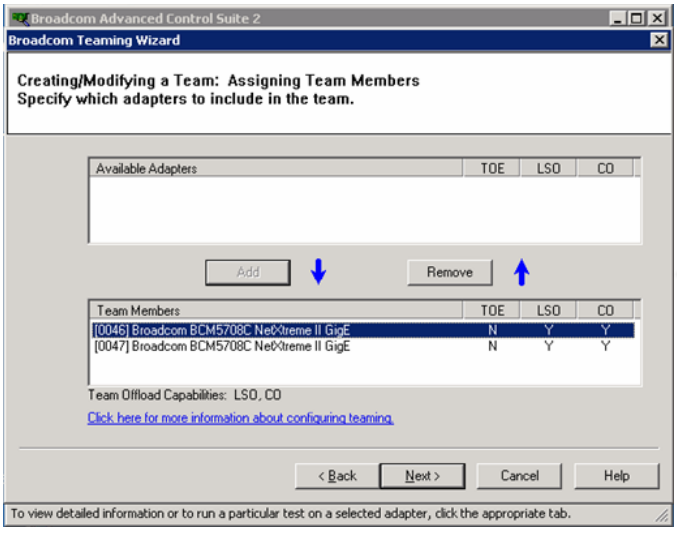

Add appropriate NICs to the newly created team

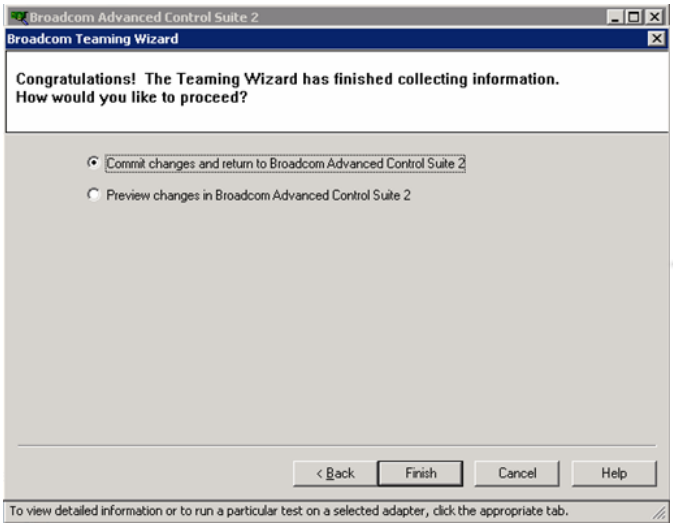

Apply changes to create the team

Configuration option to create VLAN support (802.1Q) tagging on the team

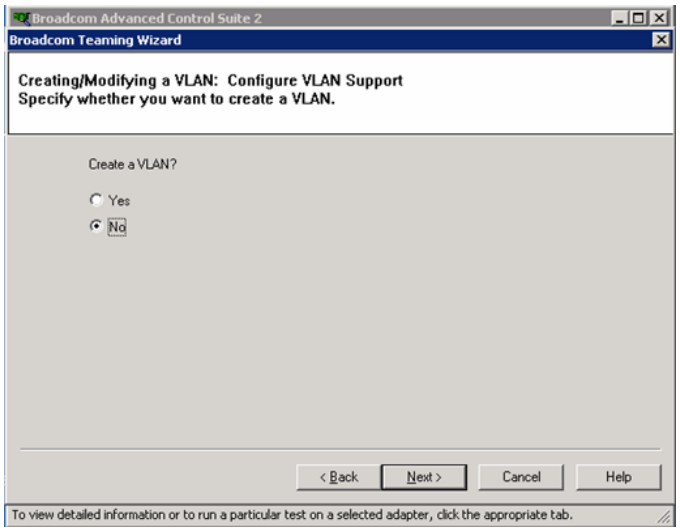

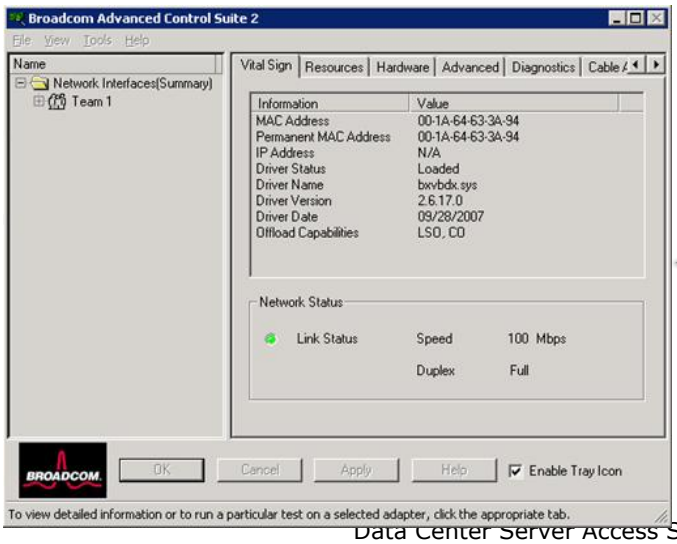

NIC team details

5. Intel NIC Teaming

Intel supports multiple different modes of NIC teaming. Each of these modes will be described briefly here, followed by the Avaya recommendations for which modes to use with various Avaya Ethernet switch solutions.

5.1 Adapter Fault Tolerance

Adapter Fault Tolerance (AFT) provides the safety of an additional backup link between the server and switch. In the case of switch port, cable, or adapter failure, network connectivity is maintained. Adapter Fault Tolerance is implemented with a primary adapter and one or more backup or secondary adapters. During normal operation, the backup adapters are in standby. If the link to the primary adapter fails, the link to the secondary adapter automatically takes over. To use Adapter Fault Tolerance all adapters must be connected to the same subnet.

5.2 Switch Fault Tolerance

Switch Fault Tolerance (SFT) teaming allows you to connect each of two teamed adapters to a separate switch. Switch Fault Tolerance can detect failures when they occur on:

- Either teamed adapter
- Either cable connecting the teamed adapter to its switch
- Switch ports connected to the adapters, if link is lost

In SFT teams, one adapter is the primary adapter and one adapter is the secondary adapter. During normal operation, the secondary adapter is in standby. In standby, the adapter is inactive and waiting for failover to occur. It does not transmit or receive other network traffic. If the primary adapter loses connectivity, the secondary adapter automatically takes over.

In SFT teams, each adapter in the team can operate at a different speed than the other.

5.3 Adaptive Load Balancing

Adaptive Load Balancing (ALB) uses software to balance routable traffic among a team of two to eight adapters (the team must include at least one server adapter) connected to the same subnet. The software analyzes the send and transmit loading on each adapter and balances the rate across the adapters based on destination address. Adapter teams configured for ALB also provide the benefits of fault tolerance.

- ALB does not load balance non-routed protocols such as NetBEUI and some IPX traffic
- You can create an ALB team with mixed speed adapters. The load is balanced according to the lowest common denominator of adapter capabilities and the bandwidth of the channel.
- On Windows systems, Receive Load Balancing is enabled by default.
- ALB is switch independent and supports switch fault tolerance by allowing the teamed ports to be connected to more than one switch in the same LAN
- It takes advantage of all available bandwidth (No standby links)
- Only use Adaptive Load Balancing when connecting to switches without any link aggregation

5.4 Static Link Aggregation

The Static Link Aggregation type of team is very similar to the Link Aggregation (802.3ad) team in that all adapters in the team need to be configured to receive packets for the same MAC address. The Static Link Aggregation team, however, does not provide LACP or marker protocol support. This type of team supports a variety of environments in which the adapter link partners are statically configured to support a proprietary trunking mechanism. For instance, this type of team could be used to support Lucent's OpenTrunk or Cisco's Fast EtherChannel (FEC) or Avaya's SMLT. Basically, this type of team is a light version of the Link Aggregation (802.3ad) type of team. This approach is much simpler, in that there is not a formalized link aggregation control protocol (LACP). As with the other types of teams, the creation of teams and the allocation of physical adapters to various teams is done statically through user configuration software. The Static Link Aggregation type of team supports load balancing and failover for both outbound and inbound traffic which is provided by the logic that is implemented on the switches.

5.5 IEEE 802.3ad Dynamic Link aggregation

IEEE 802.3ad, which is also known as Link Aggregation Control Protocol (LACP), is an IEEE specification. The main advantage of LACP is it is a standard based implementation which helps in providing interoperability with switches from different vendors.

5.6 Avaya Recommendations

The following figures and tables highlight the different Ethernet Infrastructure topologies for the Server Access Layer along with the Avaya recommended NIC Teaming mode to be used.

Multilink Trunking / Link Aggregation to Modular / Standalone / Stack

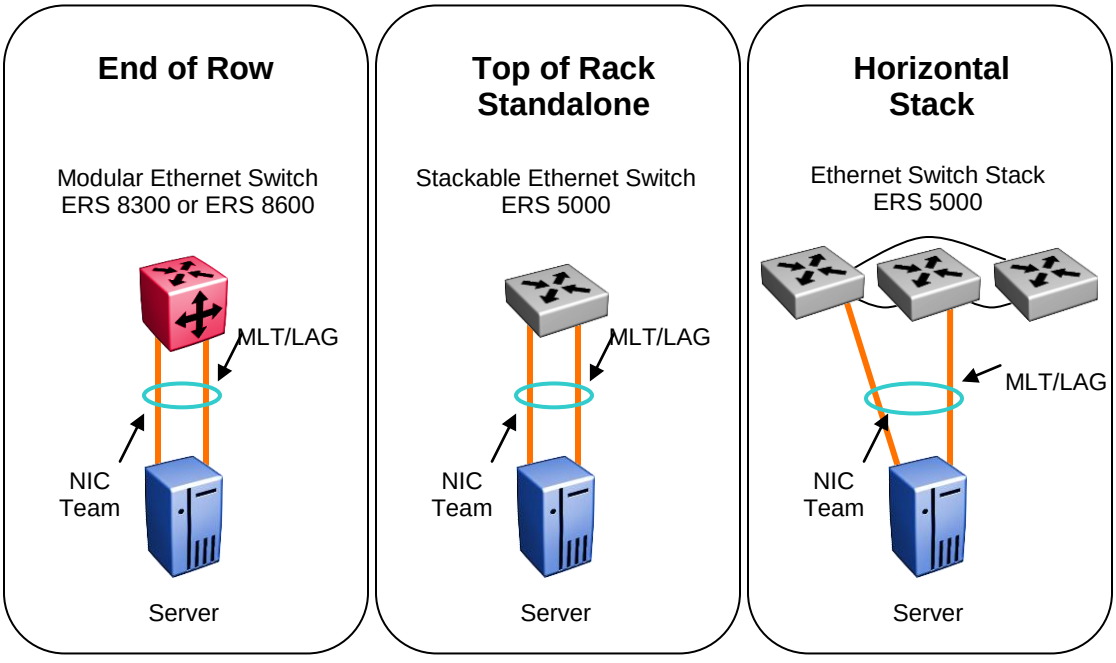

Figure 5.1: MLT/LAG with Intel NICs

Ethernet Infrastructure	ERS Platform	MLT/LAG Groups	Ports per MLT/LAG Group
Standalone Switch/Stack using MLT to Server	ERS 5000	32	Static Link Aggregation Or IEEE 802.3ad Dynamic Link Aggregation
	ERS 8300	31	
	ERS 8600	128	

Notes:

- Total number of MLT/LAG groups equals total number of NIC teamed servers supported
- Do not use Adaptive Load Balancing with this Ethernet Infrastructure topology – it is not compatible with any type of Link Aggregation techniques.

Switch Clustering to Modular / Standalone / Stack

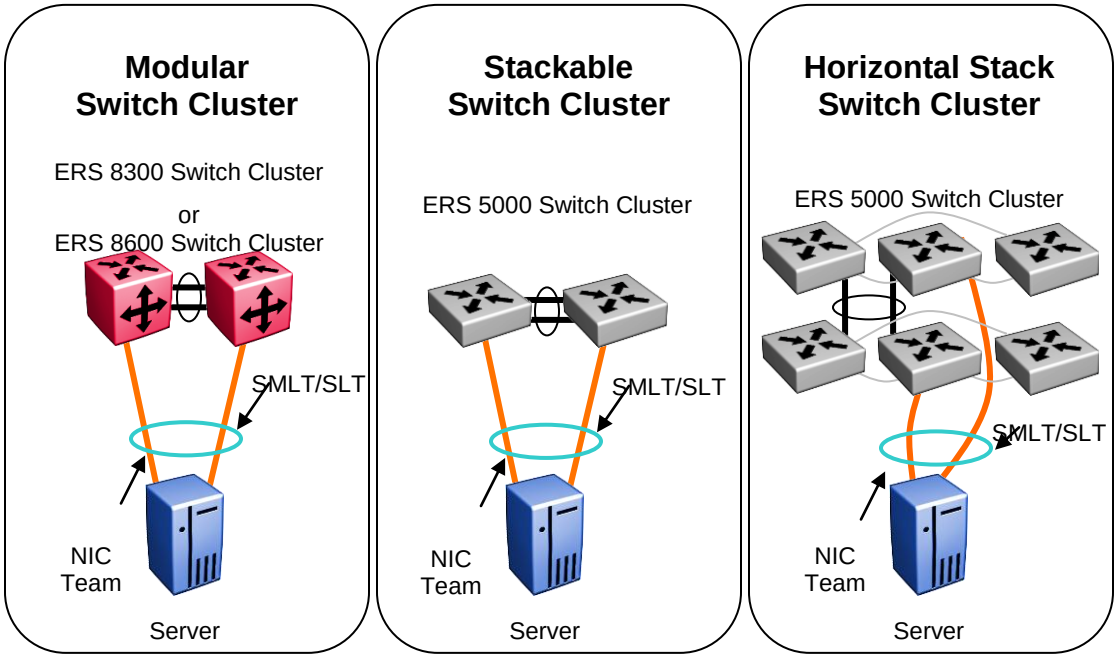

Figure 5.2: Switch Clustering with Intel NICs

Ethernet Infrastructure	ERS Platform	SMLT Groups	SLT Groups	NIC Teaming Mode
Switch Clustering to Server	ERS 5000	31	382	Static Link Aggregation
	ERS 8300	30	382	Static Link Aggregation Or IEEE 802.3ad Dynamic Link Aggregation
	ERS 8600	127	382	

Note:

- SMLT and SLT can be used on the same Switch Cluster
- If using the ERS 8600, configure LACP SMLT-SYS-ID on each of the ERS 8600s
- Do not use IEEE 802.3ad Dynamic Link Aggregation with ERS 5000 Switch Clusters as LACP over SMLT/SLT is not presently supported.
- Do not use Adaptive Load Balancing with this Ethernet Infrastructure topology – it is not compatible with any type of Link Aggregation techniques.

5.7 NIC Team Configuration Example

This section shows the details steps to configure the Intel NIC team for Static Link Aggregation – which is the preferred NIC teaming mode to work with the Ethernet infrastructure topologies discussed in the previous section.

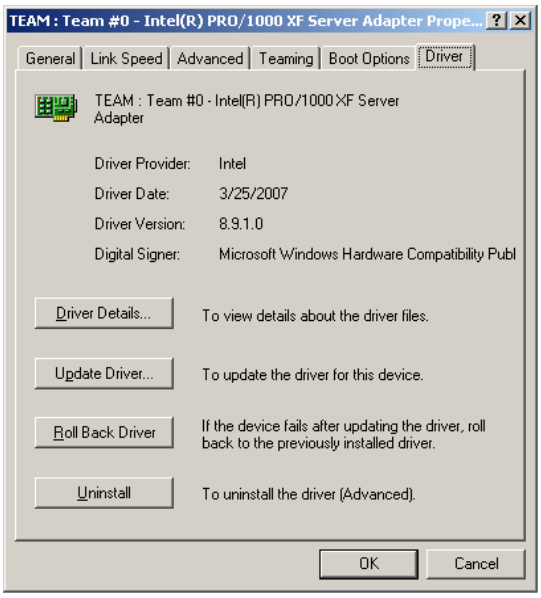

Provides the Driver details after you click on network properties tab

Start the NIC team setup by clicking the Teaming tab and click on the New Team tab

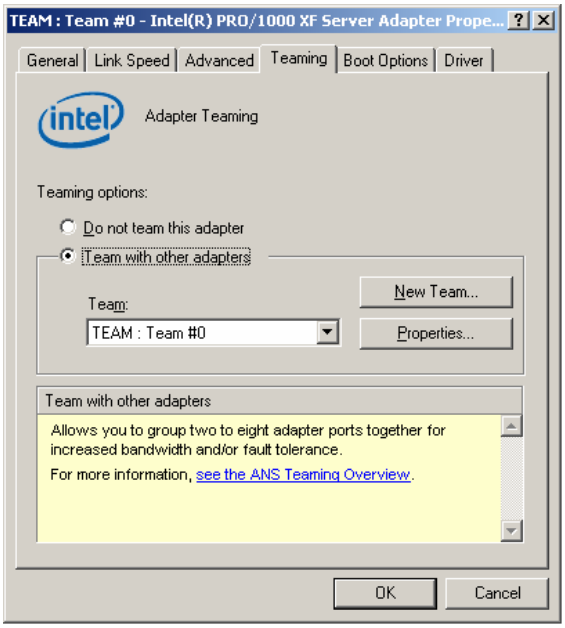

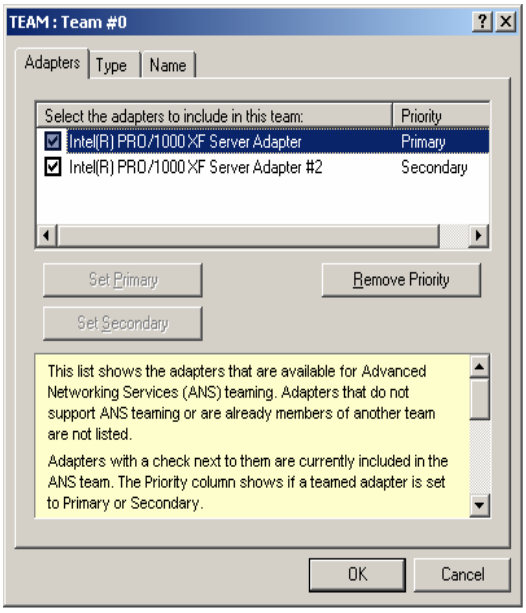

Select the adapters to be added to the team

Create NIC Team name by choosing the Name tab

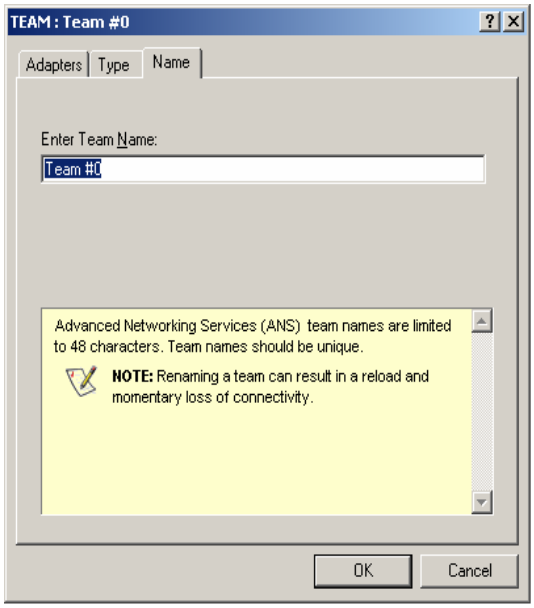

Select Type tab and choose Static Link Aggregation

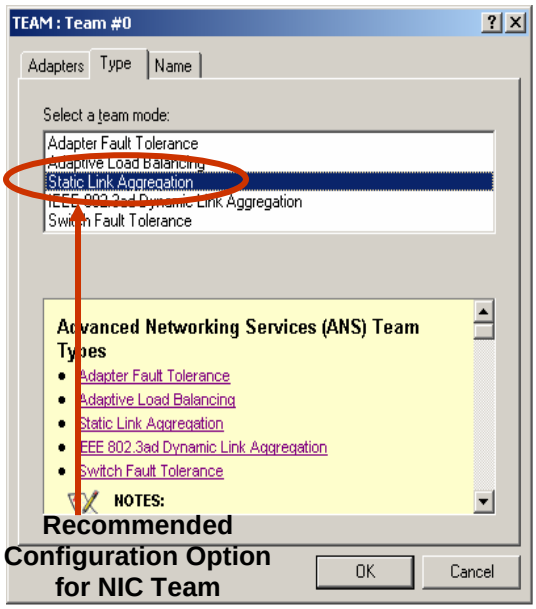

Recommended Configuration Option for NIC Team

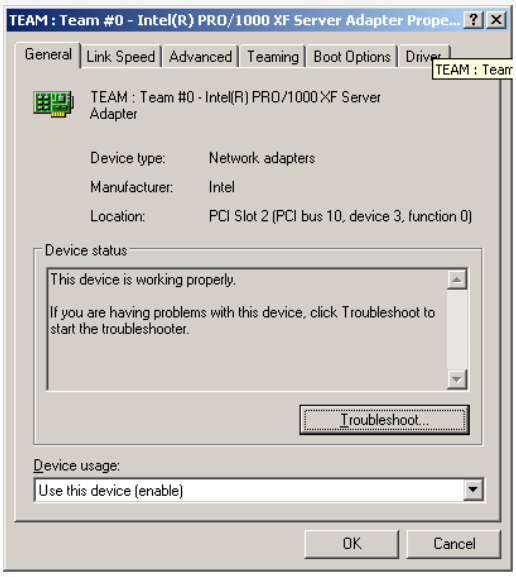

By pressing ok in the previous screen the team is configured. Select the General tab to get information on the NIC team

Network Interface Card properties

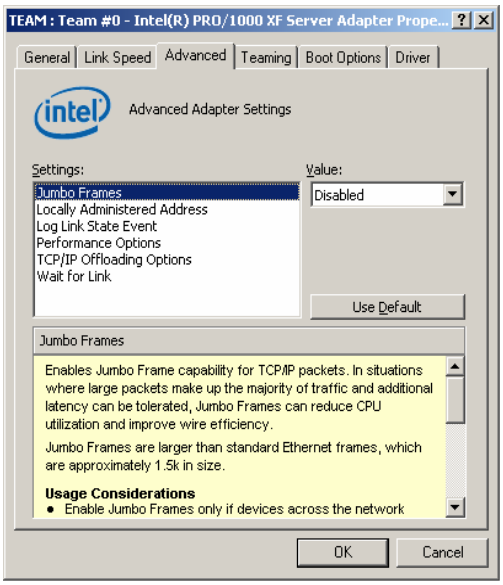

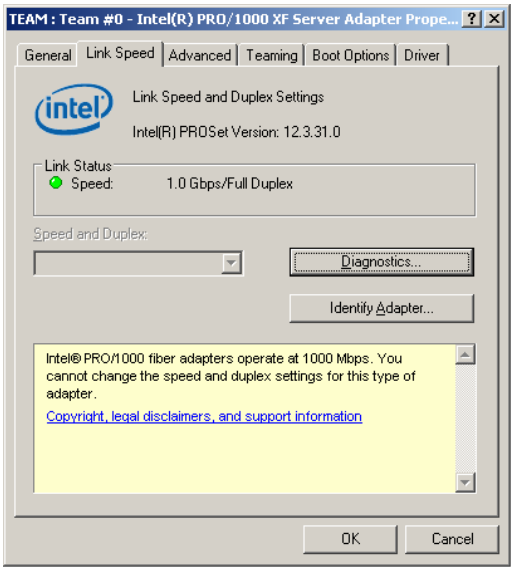

Link speed of the NIC

6. HP NIC Teaming

HP Integrity Network Adapter Teaming is a collection of fault-tolerant and load-balancing features that work together in different combinations called team types. Each of these team types will be described briefly here, followed by the Avaya recommendations for which modes to use with various Avaya Ethernet switch solutions.

6.1 Network Fault Tolerance Only (NFT)

Network Fault Tolerance (NFT) is the foundation of HP Integrity Network Adapter Teaming. In NFT mode, from two to eight ports are teamed together to operate as a single virtual network adapter. However, only one teamed port – the Primary teamed port – is used for both transmit and receive communication with the server. The remaining adapters are considered to be standby (or secondary adapters) and are referred to as Non-Primary teamed ports. Non-Primary teamed ports remain idle unless the Primary teamed port fails. All teamed ports may transmit and receive heartbeats, including Non-Primary adapters.

6.2 Network Fault Tolerance Only w/ Preference Order

Network Fault Tolerance Only with Preference Order is identical in almost every way to NFT. The only difference is that this team type allows an administrator to prioritize the order in which teamed ports should be the Primary teamed port. This ability is important in environments where one or more teamed ports are more preferred than other ports in the same team. The need for ranking certain teamed ports better than others can be a result of unequal speeds, better adapter capabilities (for example, higher receive/transmit descriptors or buffers, interrupt coalescence, etc.), or preference for the team's Primary port to be located on a specific switch.

6.3 Transmit Load Balancing w/ Fault Tolerance (TLB)

Transmit Load Balancing with Fault Tolerance (TLB) is a team type that allows the server to load balance it's transmit traffic. TLB is switch independent and supports switch fault tolerance by allowing the teamed ports to be connected to more than one switch in the same LAN. With TLB, traffic received by the server is not load balanced. The primary teamed port is responsible for receiving all traffic destined for the server. In case of a failure of the primary teamed port, the NFT mechanism ensures connectivity to the server is preserved by selecting another teamed port to assume the role.

6.4 Transmit Load Balancing w/ Fault Tolerance and Preference Order

Transmit Load Balancing with Fault Tolerance and Preference Order is identical in almost every way to TLB. The only difference is that this team type allows an administrator to prioritize the order in which teamed ports should be the Primary teamed port. This ability is important in environments where one or more teamed ports are more preferred than other ports in the same team. The need for ranking certain teamed ports better than others can be a result of unequal speeds, better adapter capabilities (for example, higher receive/transmit descriptors or buffers, interrupt coalescence, etc.), or preference for the team's Primary port to be located on a specific switch.

6.5 Switch-Assisted Load Balancing w/ Fault Tolerance (SLB)

Switch-assisted Load Balancing with Fault Tolerance (SLB) is a team type that allows full transmit and receive load balancing. SLB requires the use of a switch that supports some form of Port Trunking (for example, EtherChannel, MultiLink Trunking, etc.). SLB is similar to the 802.3ad Dynamic team type discussed later. The only difference is SLB is static configuration while 802.3ad is Dynamic configuration.

6.6 802.3ad Dynamic w/ Fault Tolerance

802.3ad Dynamic with Fault Tolerance is identical to SLB except that the switch must support the IEEE 802.3ad dynamic configuration protocol called Link Aggregation Control Protocol (LACP). In addition, the switch port to which the teamed ports are connected must have LACP enabled. The main benefit of 802.3ad Dynamic is that an Administrator will not have to manually configure the switch

6.7 Avaya Recommendations

The following figures and tables highlight the different Ethernet Infrastructure topologies for the Server Access Layer along with the Avaya recommended NIC Teaming mode to be used.

Multilink Trunking / Link Aggregation to Modular / Standalone / Stack

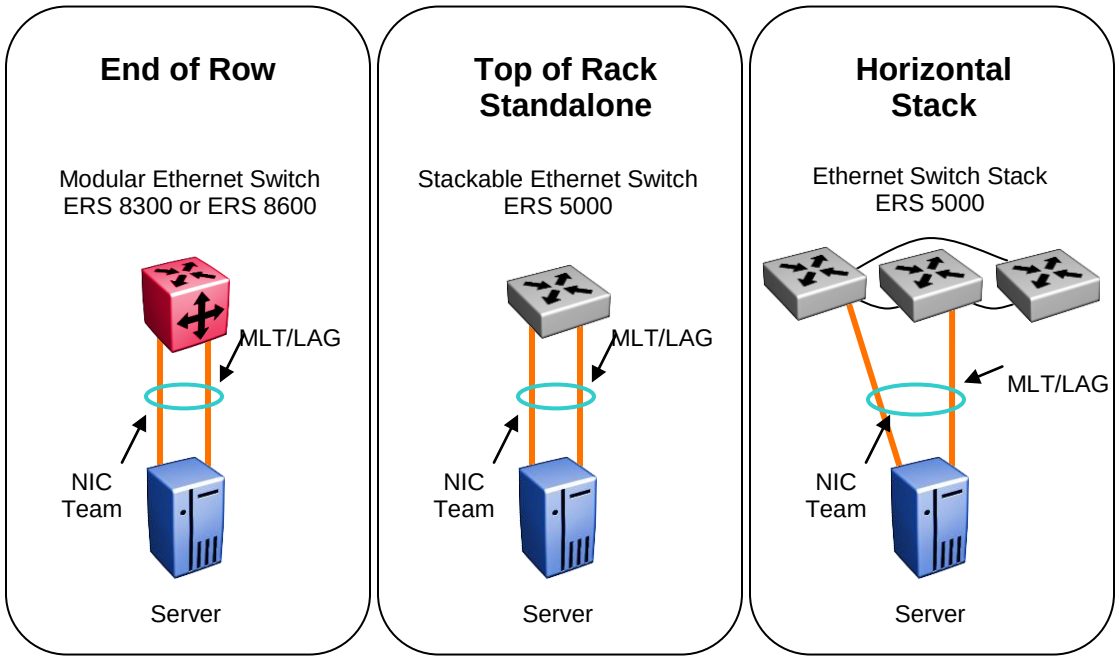

Figure 6.1: MLT/LAG with HP NICs

Ethernet Infrastructure	ERS Platform	MLT/LAG Groups	Ports per MLT/LAG Group	NIC Teaming Mode
Standalone Switch/Stack using MLT to Server	ERS 5000	32	8	Switch-assisted Load Balancing with Fault Tolerance Or 802.3ad Dynamic with Fault Tolerance
	ERS 8300	31	4	
	ERS 8600	128	8	

Notes:

- Total number of MLT/LAG groups equals total number of NIC teamed servers supported
- Do not use Transmit Load Balancing (TLB) with this Ethernet Infrastructure topology – it is not compatible with any type of Link Aggregation techniques.

Switch Clustering to Modular / Standalone / Stack

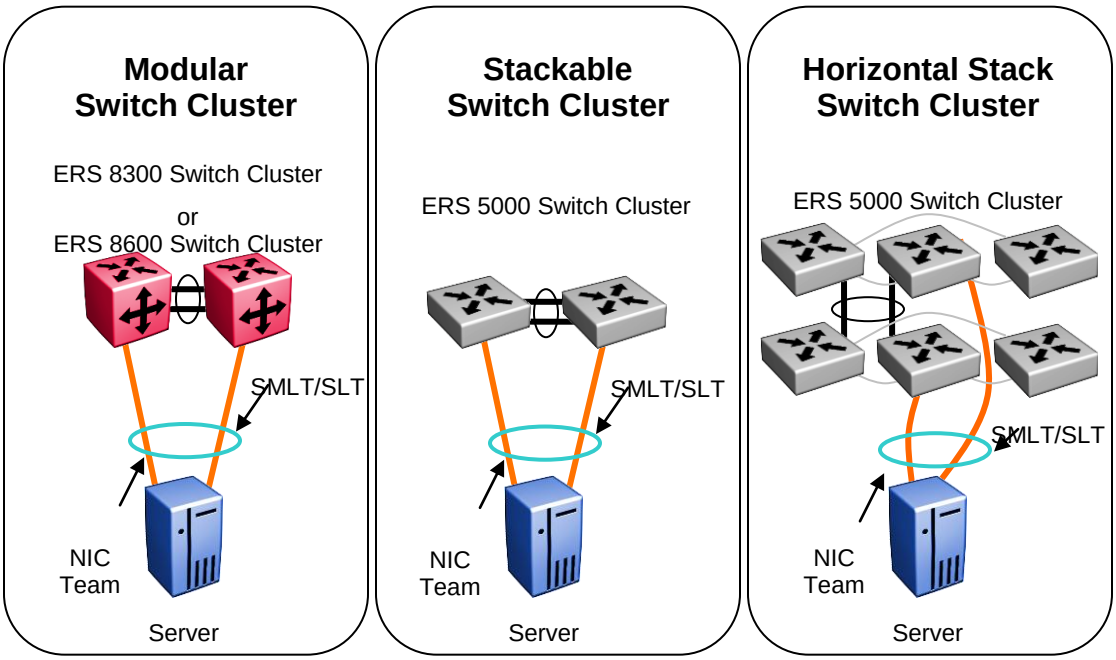

Figure 6.2: Switch Clustering with HP NICs

Ethernet Infrastructure	ERS Platform	SMLT Groups	SLT Groups	NIC Teaming Mode
Switch Clustering to Server	ERS 5000	31	382	Switch-assisted Load Balancing with Fault Tolerance
	ERS 8300	30	382	Switch-assisted Load Balancing with Fault Tolerance Or 802.3ad Dynamic with Fault Tolerance
	ERS 8600	127	382	

Note:

- SMLT and SLT can be used on the same Switch Cluster
- If using the ERS 8600, configure LACP SMLT-SYS-ID on each of the ERS 8600s
- Do not use 802.3ad Dynamic with Fault Tolerance with ERS 5000 Switch Clusters as LACP over SMLT/SLT is not presently supported.
- Do not use Transmit Load Balancing (TLB) with this Ethernet Infrastructure topology – it is not compatible with any type of Link Aggregation techniques.

6.8 NIC Team Configuration Example

This section shows the details steps to configure the HP NIC team for Switch Assisted Load Balancing with Fault Tolerance – which is the preferred NIC teaming mode to work with the Ethernet infrastructure topologies discussed in the previous section.

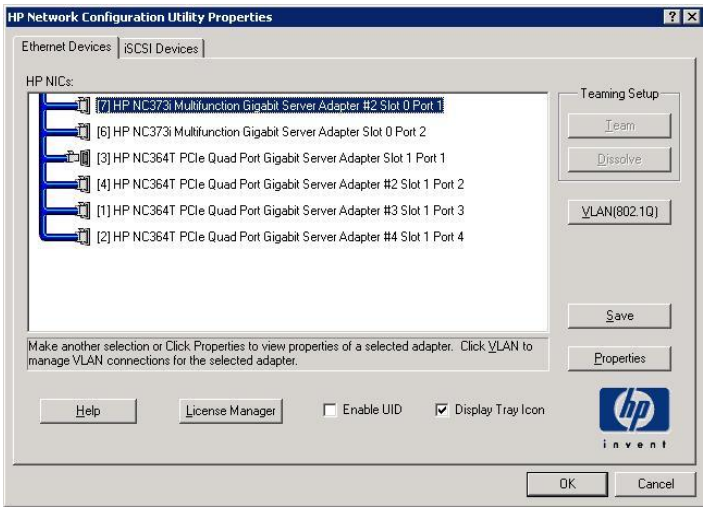

First screen after launching the HP Team Integrity software – shows the number of adapters in

Choose the adapters that you want to team and click Team tab

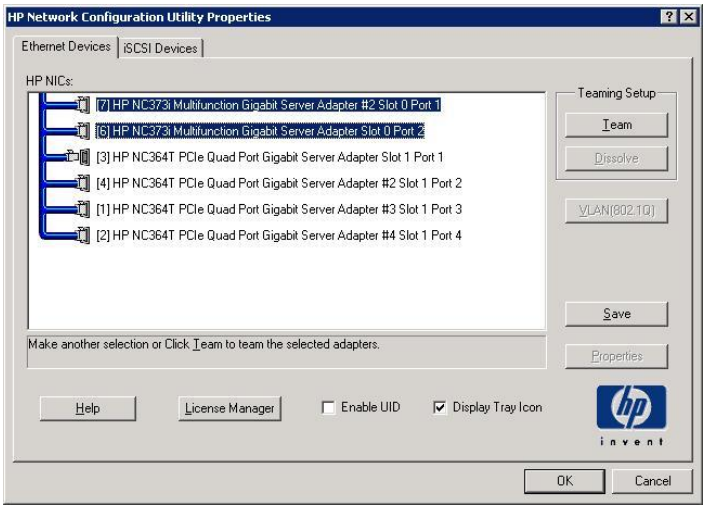

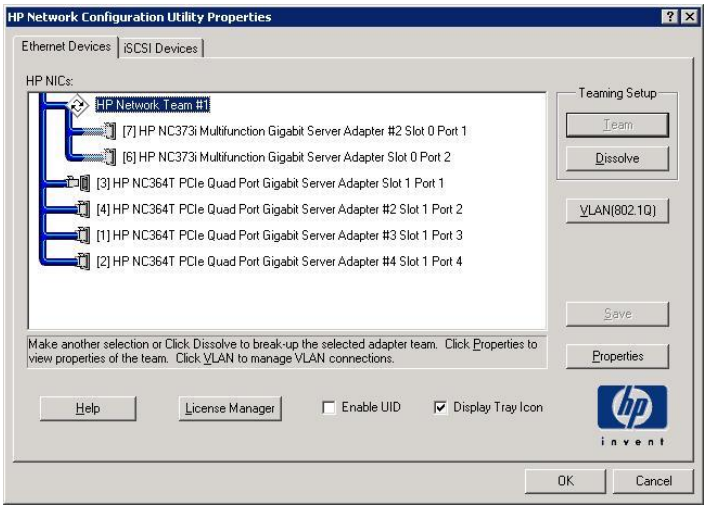

The adapters selected as part of a team in previous screen are shown under HP Network Team#1

Click on Yes tab

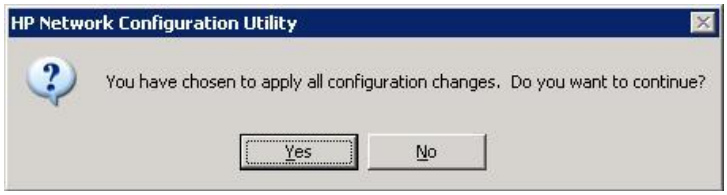

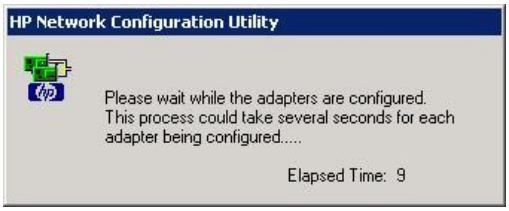

The adapters are now configured to be part of the team. Next step is to configure the team type

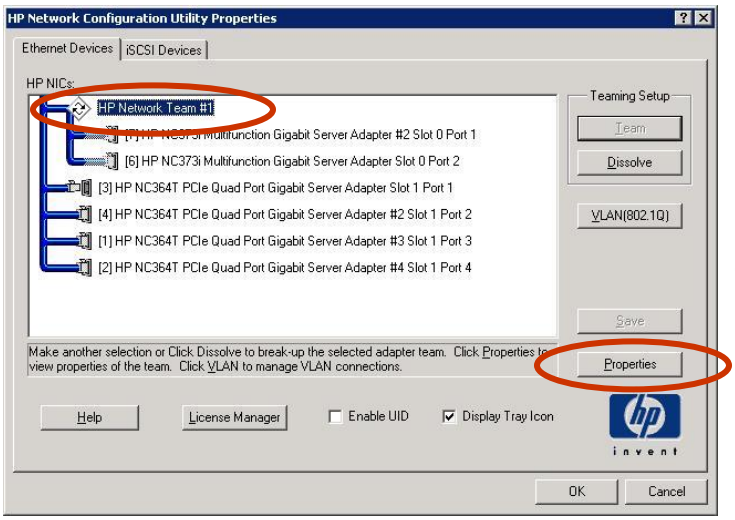

First step to configure team type is selecting the Team#1 and then click on properties

Select Switch-assisted Load Balancing with Fault Tolerance (SLB)

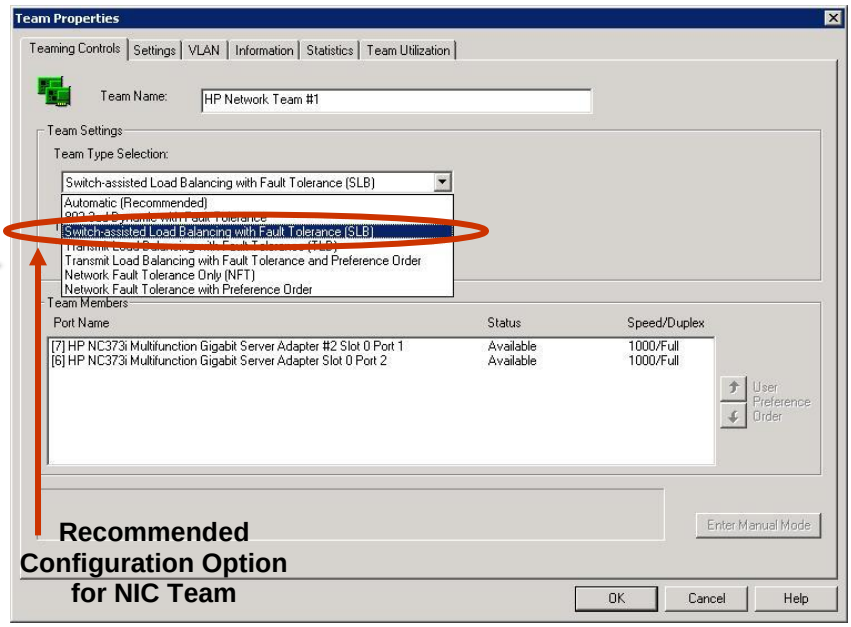

Recommended Configuration Option for NIC Team

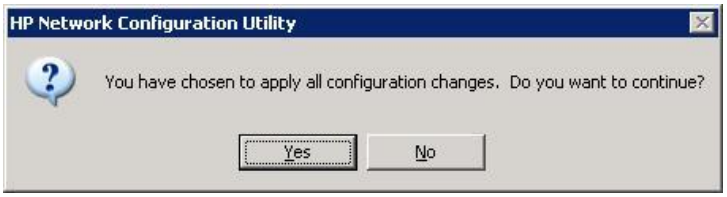

7. Avaya Ethernet Switch Configuration

This section shows the detailed steps to configure Avaya MLT and Avaya Switch Clustering to support NIC teaming.

7.1 Multilink Trunk on Avaya ERS 5650TD

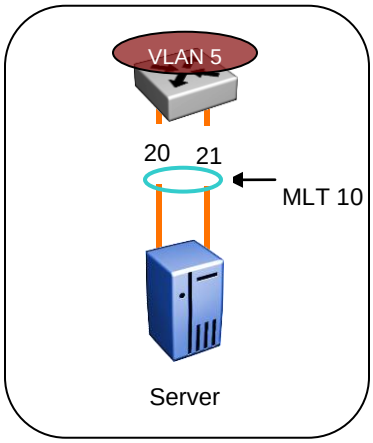

Create VLAN 5 (server) on ERS5650-1

```
5650-1(config)# vlan create 5 name server type port
5650-1(config)# vlan members remove 1 20,21
5650-1(config)# vlan members add 5 20,21
```

The above steps assume that the ERS5000 switch is using the VLAN configuration mode of *strict* (default setting). In this mode, you must first remove port members from the default VLAN 1 prior to adding these port members to a new VLAN. The VLAN configuration mode is set by using the command *vlan configcontrol <automatic|autopvid|flexible|strict>*

Create MLT 10 on ERS5650-1

```
5650-1(config)# mlt 10 member 20,21 learning disable
5650-1(config)# mlt 10 enable
```

When configuring the MLT group, it is imperative that Spanning Tree is disabled on the MLT going to the Server. The *learning disable* on the MLT command disables Spanning Tree for the MLT.

7.2 Single Link Trunking (SLT) on ERS 5650TD

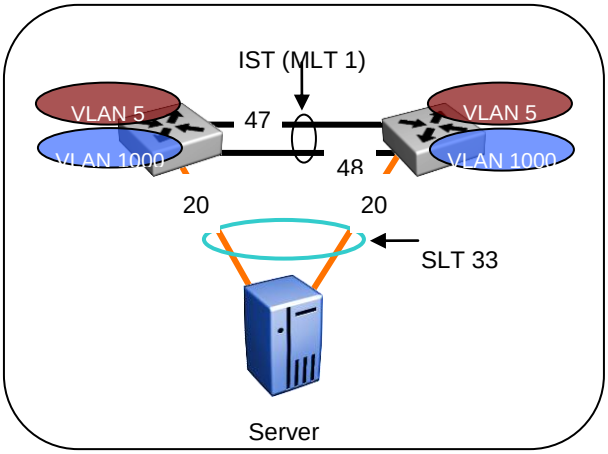

Create VLAN 5 (server) and VLAN 1000 (IST) on ERS5650-1 & ERS5650-2

```
5650-1(config)# vlan create 5 name server type port
5650-1(config)# vlan create 1000 name IST type port
5650-1(config)# vlan ports 47,48 tagging tagAll filter-untagged-frame enable
5650-1(config)# vlan members remove 1 20,47,48
5650-1(config)# vlan members add 5 20
5650-1(config)# vlan members add 1000 47,48
--------------------------------------------------------------------------------
5650-2(config)# vlan create 5 name server type port
5650-2(config)# vlan create 1000 name IST type port
5650-2(config)# vlan ports 47,48 tagging tagAll filter-untagged-frame enable
5650-2(config)# vlan members remove 1 20,47,48
5650-2(config)# vlan members add 5 20,47,48
5650-2(config)# vlan members add 1000 47,48
```

The above steps assume that the ERS5000 switch is using the VLAN configuration mode of *strict* (default setting). In this mode, you must first remove port members from the default VLAN 1 prior to adding these port members to a new VLAN. The VLAN configuration mode is set by using the command *vlan configcontrol <automatic|autopvid|flexible|strict>*

Create MLT 1 for IST

```
5650-1(config)# mlt 1 name ist enable member 47,48 learning disable
--------------------------------------------------------------------------------
5650-2(config)# mlt 1 name ist enable member 47,48 learning disable
```

When configuring the MLT group, it is imperative that Spanning Tree is disabled on the MLT going to the Server. The *learning disable* on the MLT command disables Spanning Tree for the MLT.

Enable IP Routing Globally and add IP address to the IST VLAN

```
5650-1(config)# ip routing
5650-1(config)# interface vlan 1000
5650-1(config-if)# ip address 10.1.2.1 255.255.255.252
5650-1(config-if)# exit
--------------------------------------------------------------------------------
5520-1(config)# ip routing
5520-1(config)# interface vlan 1000
5520-1(config-if)# ip address 10.1.2.2 255.255.255.252
5520-1(config-if)# exit
```

Create IST

```
5650-1(config)# interface mlt 1
5650-1(config-if)# ist enable peer-ip 10.1.2.2 vlan 1000
5650-1(config-if)# exit
--------------------------------------------------------------------------------
5650-2(config)# interface mlt 1
5650-2(config-if)# ist enable peer-ip 10.1.2.1 vlan 1000
5650-2(config-if)# exit
```

Configure the VLACP MAC and enable VLACP globally

```
5650-1(config)# vlacp macaddress 180.c200.f
5650-1(config)# vlacp enable
--------------------------------------------------------------------------------
5650-2(config)# vlacp macaddress 180.c200.f
5650-2(config)# vlacp enable
```

It is recommended to use the reserved multicast MAC address of 01:80:c2:00:00:0f for the VLACP MAC address. On the ERS5000, enter the hex value *180.c200.f*.

Enable VLACP on IST

```
5650-1(config)# interface FastEthernet all
5520-1(config-if)# vlacp port 47,48 timeout long
5650-1(config-if)# vlacp port 47,48 slow-periodic-time 10000
5650-1(config-if)# vlacp port 47,48 timeout-scale 3
5650-1(config-if)# vlacp port 47,48 enable
5650-1(config-if)# exit
--------------------------------------------------------------------------------
5650-2(config)# interface FastEthernet all
5520-2(config-if)# vlacp port 47,48 timeout long
5650-2(config-if)# vlacp port 47,48 slow-periodic-time 10000
5650-2(config-if)# vlacp port 47,48 timeout-scale 3
5650-2(config-if)# vlacp port 47,48 enable
5650-2(config-if)# exit
```

Create SLT to Server

```
5520-1(config)# interface FastEthernet ALL
5520-1(config-if)# smlt port 20 33
5520-1(config-if)# exit
--------------------------------------------------------------------------------
5520-1(config)# interface FastEthernet ALL
5520-1(config-if)# smlt port 20 33
5520-1(config-if)# exit
```

Spanning Tree must be disabled on the SLT ports. When the Ethernet ports were removed from VLAN 1 above, they were also removed from the Spanning Tree group. When they were added to VLAN 5, they will be put back into the Spanning Tree group associated with VLAN 5, however, their Spanning Tree State will remain disabled.

In the case where Spanning Tree is enabled on the SLT port, use the following commands on each switch to disable:

5650-1(config)# ***interface fastEthernet <port #>***

5650-1(config-if)# ***spanning-tree learning disable***

7.3 Split Multilink Trunking (SMLT) on ERS8600

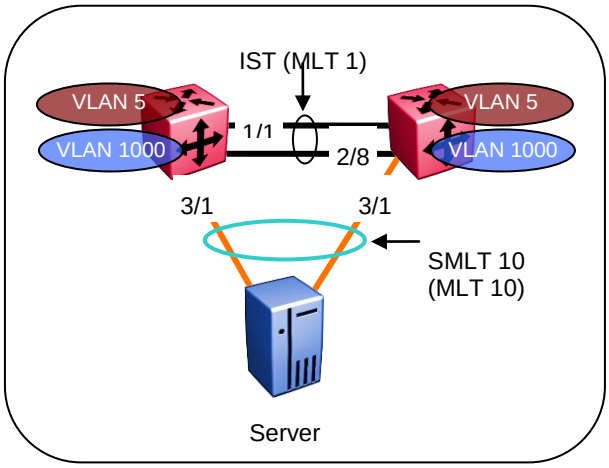

For this configuration example, ERS8600-1 is configured using the ACLI command interface while ERS8600-2 is configured using the Passport command interface.

Create VLANs 5 and 1000

```
ERS8600-1:5(config)# vlan create 1000 name IST type port 1
ERS8600-1:5(config)# interface vlan 1000
ERS8600-1:5(config-if)# ip address 10.1.2.1 255.255.255.252
ERS8600-1:5(config-if)# exit
ERS8600-1:5(config)# vlan create 5 name Server type port 1
ERS8600-1:5(config)# interface vlan 5
ERS8600-1:5(config-if)# ip address 10.1.100.2 255.255.255.0
ERS8600-1:5(config-if)# exit
```

```
ERS8600-2:5# config vlan 1000 create byport 1 name IST
ERS8600-2:5# config vlan 1000 ip create 10.1.2.2/30
ERS8600-2:5# config vlan 5 create byport 1 name Server
ERS8600-2:5# config vlan 5 ip create 10.1.100.3/24
```

Create MLT 1 for IST

```
ERS8600-1:5(config)# mlt 1
ERS8600-1:5(config)# mlt 1 name IST
ERS8600-1:5(config)# mlt 1 member 1/1,2/8
```

```
ERS8600-1:5(config)# mlt 1 encapsulation dot1q
ERS8600-1:5(config)# vlan mlt 5 1
```

```
ERS8600-2:5# config mlt 1 create
ERS8600-2:5# config mlt 1 name IST
ERS8600-2:5# config mlt 1 add port 1/1,2/8
ERS8600-2:5# config vlan 5 add-mlt 1
```

Create IST

```
ERS8600-1:5(config)# interface mlt 1
ERS8600-1:5(config-mlt)# ist peer-ip 10.1.2.2 vlan 1000
ERS8600-1:5(config-mlt)# ist enable
ERS8600-1:5(config-mlt)# end
```

```
ERS8600-2:5# config mlt 1 ist create ip 10.1.2.1 vlan-id 1000
ERS8600-2:5# config mlt 1 ist enable
```

Enable VLACP on IST

```
ERS8600-1:5(config)# interface gigabitEthernet 1/1,2/8
ERS8600-1:5(config-if)# vlacp slow-periodic-time 10000
ERS8600-1:5(config-if)# vlacp funcmac-addr 01:80:c2:00:00:0f
ERS8600-1:5(config-if)# vlacp enable
ERS8600-1:5(config-if)# exit
```

```
ERS8600-2:5# config ethernet 1/1,2/8 vlacp slow-periodic-time 10000
ERS8600-2:5# config ethernet 1/1,2/8 vlacp macaddress 01:80:c2:00:00:0f
ERS8600-2:5# config ethernet 1/1,2/8 vlacp enable
```

Create SMLT to Server

```
ERS8600-1:5(config)# mlt 10
ERS8600-1:5(config)# mlt 10 member 3/1 vlan 5
ERS8600-1:5(config)# mlt 10 encapsulation dot1q
ERS8600-1:5(config)# interface mlt 10
ERS8600-1:5(config-mlt)# smlt 10
ERS8600-1:5(config-mlt)# end
```

```
ERS8600-2:5# config mlt 10 create
ERS8600-2:5# config mlt 10 perform-tagging enable
ERS8600-2:5# config mlt 10 add port 3/1
ERS8600-2:5# config vlan 5 add-mlt 10
ERS8600-2:5# config mlt 10 smlt create smlt-id 10
```

Spanning Tree is automatically disabled when creating the SMLT ports on the ERS 8600.

7.3.1 Optional Configuration – Default Gateway Redundancy

If the ERS 8600 Switch Cluster is running Layer 3 default gateway redundancy can be configured using either VRRP or RSMLT Layer 2 Edge. Do not configure both VRRP and RSMLT Layer 2 Edge on the same VLAN.

7.3.2 VRRP Configuration

Create VRRP VIP

```
ERS8600-1:5(config)# interface vlan 5
ERS8600-1:5(config-if)# ip vrrp address 5 10.1.100.1
```

```
ERS8600-2:5# config vlan 5 ip vrrp 5 address 10.1.100.1
```

Enable Backup Master

```
ERS8600-1:5(config-if)# ip vrrp 5 backup-master enable
```

```
ERS8600-2:5# config vlan 5 ip vrrp 5 backup-master enable
```

Set the hold down timer to 90 seconds

```
ERS8600-1:5(config-if)# ip vrrp 5 holddown-timer 90
```

```
ERS8600-2:5# config vlan 5 ip vrrp 5 holddown-timer 90
```

Set VRRP priority (only need to change priority on one switch)

```
ERS8600-1:5(config-if)# ip vrrp 5 priority 110
```

Enable VRRP

```
ERS8600-1:5(config-if)# ip vrrp 5 enable
ERS8600-1:5(config-if)# exit
```

```
ERS8600-2:5# config vlan 5 ip vrrp 5 enable
```

7.3.2.1 RSMLT Layer 2 Edge Configuration

Enable RSMLT
ERS8600-1:5(config)# ***interface vlan 5***
ERS8600-1:5(config-if)# ***ip rsmlt***
ERS8600-2:5# ***config vlan 5 ip rsmlt enable***
Set the RSMLT hold-up timer to infinity
ERS8600-1:5(config-if)# ***ip rsmlt holdup-timer 9999***
ERS8600-1:5(config-if)# ***exit***
ERS8600-2:5# ***config vlan 5 ip rsmlt holdup-timer 9999***
Enable RSMLT-edge support
ERS8600-1:5(config)# ***ip rsmlt edge-support***
ERS8600-2:5# ***config ip rsmlt rsmlt-edge-support enable***

7.3.3 Optional Configuration – LACP between Server and Switch Cluster

If necessary, the ERS 8600 and ERS 8300 can be configured for LACP over SMLT to the server. Use this configuration only when configuring the NIC team for 802.3ad Link Aggregation. The following configuration would replace the SMLT to Server configuration shown earlier.

Create SMLT to Server
ERS8600-1:5(config)# ***lacp enable***
ERS8600-1:5(config)# ***lacp smlt-sys-id 00:e0:7b:82:9c:00***
ERS8600-1:5(config)# ***interface gigabitethernet 3/1***
ERS8600-1:5(config-if)# ***lacp key 1***
ERS8600-1:5(config-if)# ***lacp enable***
ERS8600-1:5(config-if)# ***exit***
ERS8600-1:5(config)# ***vlan members remove 1 3/1***
ERS8600-1:5(config)# ***vlan members add 5 3/1***
ERS8600-1:5(config)# ***mlt 10***
ERS8600-1:5(config)# ***mlt 10 encapsulation dot1q***
ERS8600-1:5(config)# ***interface mlt 10***
ERS8600-1:5(config-mlt)# ***lacp enable key 1***
ERS8600-1:5(config-mlt)# ***smlt 10***
ERS8600-1:5(config-mlt)# ***exit***
ERS8600-1:5(config)# ***vlan mlt 5 1***

```
ERS8600-2:5# config mlt 10 create
ERS8600-2:5# config mlt 10 perform-tagging enable
ERS8600-2:5# config mlt 10 lacp key 1
ERS8600-2:5# config mlt 10 lacp enable
ERS8600-2:5# config vlan 1 ports remove 3/1
ERS8600-2:5# config vlan 5 ports add 3/1
ERS8600-2:5# config ethernet 3/1 lacp key 1
ERS8600-2:5# config ethernet 3/1 lacp aggregation true
ERS8600-2:5# config ethernet 3/1 lacp enable
ERS8600-2:5# config mlt 10 smlt create smlt-id 10
ERS8600-2:5# config lacp smlt-sys-id 00:e0:7b:82:9c:00
ERS8600-2:5# config vlan 5 add-mlt 10
ERS8600-2:5# config lacp enable
```

Spanning Tree is automatically disabled when creating the SMLT ports on the ERS 8600.

The smlt-sys-id is the base mac address of one of the ERS 8600 switches in this example and must be configured the same on both switches.

8. 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.

8.1 Getting technical documentation

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

8.2 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.

8.3 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.

8.4 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.