



Resilient Data Center Solutions for VMware ESX Servers Technical Configuration Guide

Avaya Data Solutions

Document Date: July 2010

Document Number: NN48500-542

Document Version: 2.1

© 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/](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

This Technical Configuration Guide describes options and best practices for connecting application servers in a data center environment that will run VMware ESX Server. The guide provides Avaya recommendations and configuration examples for a resilient Ethernet switching solution that maximizes performance, resource utilization and minimizes both planned and unplanned downtime.

Table of Contents

Figures	5
Document Updates	6
Conventions	6
1. Overview: Avaya ERS Resilient Ethernet Switching and VMware ESX Server	7
2. Recommended Solution: Horizontal Stacking with Switch Clustering	9
2.1 Horizontal Stacking	10
2.2 Hybrid Stacks	11
2.3 Switch Clustering	11
2.4 Horizontal Stacking with Switch Clustering.....	13
2.5 Server Default Gateway Redundancy.....	14
2.6 Assumptions.....	16
2.7 Objectives.....	16
2.8 Recommended Connectivity	17
2.9 Configuration Steps:.....	18
2.10 Verification.....	61
3. Customer service	62
3.1 Getting technical documentation.....	62
3.2 Getting product training	62
3.3 Getting help from a distributor or reseller.....	62
3.4 Getting technical support from the Avaya Web site	62

Figures

Figure 1 – Horizontal Stacking with Switch Clustering Architecture	9
Figure 2 – ERS 5000 Horizontal Stacking Architecture	10
Figure 3 – SLT and SMLT Terminology	12
Figure 4 – Horizontal Stacking with Switch Clustering	13
Figure 4 – Horizontal Stacking with Switch Clustering	13
Figure 5 – ERS 8600 RSMLT Layer 2 Edge.....	15
Figure 6 – ERS 8600 VRRP.....	16
Figure 7 - Connectivity for VMware ESX Servers	17
Figure 8 – VMware Server Connectivity Reference Topology.....	18
Figure 9 – Layer 3 Configuration Details using RSMLT for VMware Topology	19
Figure 10 – Layer 3 Configuration Details using VRRP for VMware Topology	20

Document Updates

July 19, 2010

Conventions

This section describes the text, image, and command conventions used in this document.

Symbols:



Tip – Highlights a configuration or technical tip.



Note – Highlights important information to the reader.



Warning – Highlights important information about an action that may result in equipment damage, configuration or data loss.

Text:

Text in a Courier New font indicates text the user must enter or select in a menu item, button or command:

```
ERS5520-48T# show running-config
```

1. Overview: Avaya ERS Resilient Ethernet Switching and VMware ESX Server

As customers realize the benefits of machine virtualization in the data center, it is evident that optimizing storage communication and minimizing server and network downtime are key requirements. As an individual physical server can have many virtual machines, running critical applications, both the server hardware and the networking topology must eliminate any single points of failure, provide instant failover capability for link or node failures, and ideally maximize the use of deployed switches, network interface cards (NICs), and cabling. Furthermore, it is imperative that both planned and unplanned downtime be minimized; specifically the data center operations staff should be able to perform any maintenance tasks on the network and servers without impacting service. Avaya Switch Clustering is a solution provided on a number of Ethernet Routing Switches including the ERS 8600, ERS 8300, ERS 1600 and ERS 5000 that when combined with VMware best practices for virtual networking deliver unmatched resiliency, simplicity, performance and cost effectiveness.

This TCG will focus on resilient server access using ERS 5000 series switches deployed in horizontal stacks. Further it will illustrate the inclusion of the ERS 8600 Switch Cluster core for the routing function in an overall highly resilient square SMLT topology with use of VRRP, as well as an ultra resilient RSMLT configuration.

VMware virtual networking recommendations include:

- Separate Service Console/VMotion/iSCSI from VM Traffic
- Cross-Team On-board NICs with PCI NICs for added hardware level redundancy
- Disable Spanning Tree Protocol
- Use Virtual Switch Tagging for VLAN assignment and isolation
 - Virtual Machines are connected to a port group on the ESX virtual switch and the virtual switch applies 802.1Q VLAN tags to Ethernet frames for the port group. The SMLT links to the ERS 5000s are configured as tagged trunks to interconnect the VLANs into the rest of the Data Center network. The use of VLANs is recommended to provide flexible network configuration and partitioning. The use of VST is recommended by VMware as the preferred method of using VLANs within ESX Server.

Avaya recommendations for connecting VMware ESX include:

- Use ERS Switch Clustering (SMLT) to the server NIC level
 - Avaya Switch Clustering using Split Multilink Trunking provides added resiliency to the server access switching solution. The NIC teams from the ESX server are connected to independent ERS 5000 switches providing multi-switch resiliency. By adding Switch Clustering, we now have two fully independent switches providing forwarding. Switch Clustering extends the resiliency to include maintenance operations such as software upgrades. One of the members of the cluster could be upgraded and the switch restarted without any impact to the ESX server traffic. Traffic would continue to run over the NIC team links running to the other switch cluster member.
- Use Horizontal Stacking and Switch Clustering to simplify cabling and provide zero service impact maintenance
- Use Layer 2 at the Server Access Layer and Layer 3 at the Data Center Core/Distribution

- Use ESX NIC Teams in IP Hashing Mode to balance traffic across NIC team and take advantage of multiple switch redundancy
 - VMware “IP hash” mode NIC-Teaming spreads the traffic across all the available links in the trunk group. A given server will be serving multiple clients with different IP addresses such that a hash based on source and destination IP addresses will effectively spread the traffic. This is preferred to the MAC or Virtual NIC port teaming modes that would limit the traffic for any VM to a single NIC and link. This adds another layer of resiliency to the solution and is especially useful for iSCSI traffic to provide higher throughput than can be supported on a single NIC.

Avaya Resilient Stacking and Switch Clustering with Split Multi-Link Trunking are key concepts for the solutions discussed in this guide. For background information on these please see the Large Campus Technical Solution Guide (NN48500-575) on the Avaya Support Web site: <http://www.avaya.com/support>.

2. Recommended Solution: Horizontal Stacking with Switch Clustering

This solution provides high-density Gigabit Ethernet edge switching for VMware servers and provides multi-switch resiliency. The highlighted benefits of this solution include:

- Near Zero Down Time Upgrade: All the traffic can be directed through one cluster while the other is being upgraded with minimal interruption of traffic flows.
- Low Latency Communication: The unique horizontal stacking capability allows very low latency communication amongst devices connected anywhere within the stack. For most communications there is no impact on uplink capacity. This is ideal for rapid communication required for 'server clustering', as well as for iSCSI.
- High Fault Tolerance: Use of MLT, SMLT, and RSMLT redundancy mechanisms allow for negligible latency impact during link failures. Redundancy protocols with significant convergence times such as Spanning Tree and VRRP are no longer required.
- Physical Connection Efficiency: The stacking capability allows for high port density, while requiring only a few high capacity uplinks to the core.
- Flexible Growth: As the datacenter grows, additional switches, across multiple data cabinets can be added to the stack. New units can be added to the stack without interrupting operations of the current stack.




Figure 1 – Horizontal Stacking with Switch Clustering Architecture

2.1 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 2 – ERS 5000 Horizontal Stacking Architecture

Horizontal Stacking provides the following benefits:

- Fault-tolerant or Load sharing NIC teaming into stack
- High bandwidth and low latency between servers (9μs)
- Highly resilient stacking technology with scalable uplinks
- Flexibility to spread across multiple data cabinets (100s of servers)
- Ideal for Grid Computing / High-Performance Computing Solutions

2.2 Hybrid Stacks

Hybrid stacks of ERS 5500 and ERS 5600 switches are supported. It is important to understand that the bandwidth of the stack connections are based on each switches capabilities. In a hybrid, the stack bandwidth between the ERS 5500 and ERS 5600 is 80Gbps, while the stack bandwidth between two ERS 5600s is 144Gbps. This unique capability allows the maximum bandwidth for a stack and does not reduce the stack to the lowest common denominator.

Hardware differences between the various models within the ERS 5000 family necessitate different software functionality and scalability. The following highlights those differences when using the different switches to create a mixed or hybrid stack. A mixed stack is defined as different models within the same switch type (stacking ERS 5510s with ERS 5530s) while a hybrid stack refers to stacking of ERS 5500s with ERS 5600s.

Mixed or Hybrid stack differences when using the ERS 5510

- Filter untagged frames and VLACP cannot be used simultaneously on the ERS 5510 but can be used on non-5510 units in the stack
- Dual agent is not supported on the ERS 5510

When creating a hybrid stack of ERS 5600s and ERS 5500s

- ERS 5600 must be the base unit – in a scenario where the base unit fails, the temporary base unit (normally unit #2) must be an ERS 5600 if one exists in the stack. At no time, can an ERS 5500 unit be the base unit if an ERS 5600 unit is present in the stack
- When stacking ERS 5600 to ERS 5600, the ERS 5600 stacking cable must be used
- When stacking ERS 5600 to ERS 5500, either the ERS 5600 or ERS 5500 stacking cable can be used
- Keep ERS 5600s adjacent to each other and ERS 5500s adjacent to each other in order to maximize stacking bandwidth
- Configure the **stack oper-mode hybrid** – this command is available only on the ERS 5600 switches
- When stacking the ERS 5632FD with any ERS 5500, the last two 10 Gigabit ports on all ERS 5632FDs in the stack will be disabled

Functionality supported only on certain ERS 5000 switches

- DoS attack prevention package supported only on ERS 5600 series
- Many to Many port mirroring supported only on ERS 5600 series
- Real Time Clock supported on ERS 5530 and ERS 5600 series

2.3 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.

2.3.1 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 – SLT and SMLT Terminology

Table 1 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	67
ERS 5000	8	32	31	31	31	398	190	62

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

Table 1 – MLT/SMLT/SLT Scaling Capabilities

2.4 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 if 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 4 – Horizontal Stacking with Switch Clustering

The stacked Switch Cluster is comprised of a stack of ERS 5000 switches for each IST peer. This configuration can terminate server connections using MLT-based SMLT or port-based SLT. Design considerations for the stacked Switch Cluster include:

- Stack on each side
 - Up to 8 units per stack – mixed stacks are supported for the Switch Cluster Core, however, Avaya recommends using pure stacks of the same hardware model when possible. This eliminates any issues with reduced functionality due to hardware limitations. The ERS 5510 is limited in functionality and scalability as compared to the ERS 5520 and ERS 5530. Likewise, the ERS 55xx switches are limited in functionality and scalability as compared to the ERS 56xx switches. Mixed stacks do provide increased flexibility but at a possible cost of overall functionality and scalability.
 - Switches terminating IST should be identical hardware type within the stack and on each side of the IST link.
 - At least one of the IST ports must be on the base unit of the stack.
 - When using a stack of two switches, forced stack mode should be enabled (requires Release 6.1 software or later on the ERS 5000 platform)
- IST is two port DMLT (minimum) and up to eight port DMLT (maximum)
- SMLT or SLT supported for the Edge connections
 - SMLT requires at least two ports on each IST peer (minimum four ports required for the edge connection)
- Topology is supported with ERS 5000 version 5.1 software or later
- Two Advanced Software Licenses required (one for each IST peer stack)

2.5 Server Default Gateway Redundancy

Avaya offers two different alternatives for server default gateway redundancy. Both VRRP (Virtual Router Redundancy Protocol) and RSMLT Layer 3 Edge can provide this functionality. Each of these features are described below and configuration examples are shown for each later in this document.

2.5.1 RSMLT Layer 2 Edge

RSMLT L2 Edge offers an alternative to VRRP for server default gateway redundancy. VRRP and RSMLT L2 Edge can be used on the same Switch Cluster on different VLANs, but do not use both VRRP and RSMLT L2 Edge on the same VLAN simultaneously.

The RSMLT implementation does not use a Virtual IP address but instead uses physical IP addresses on each ERS 8600 for redundancy. RSMLT L2 Edge stores the RSMLT peer MAC/IP address-pair in its local configuration file and restores the configuration if the peer does not restore after a simultaneous reboot of both RSMLT peer switches. It is imperative to save the configuration file on each ERS 8600 when RSMLT L2 Edge is first implemented to ensure that the peer MAC/IP address-pair is saved in the configuration file.

Each ERS 8600 is able to forward on behalf of itself as well as its peer in the Switch Cluster. This makes very efficient use of bandwidth and resources and also ensures seamless failover and recovery in the event of a failure.

Avaya recommends using RSMLT L2 Edge in place of VRRP as it provides several advantages, including:

- RSMLT is only limited by the number of IP interfaces on the ERS 8600
 - VRRP is limited to 250 instances
- RSMLT requires significantly less control traffic
- RSMLT is much less intensive on CPU resources




Figure 5 – ERS 8600 RSMLT Layer 2 Edge

When implementing RSMLT L2 Edge, make sure to:

- Enable RSMLT on each VLAN
- Enable RSMLT-Edge-Support on each VLAN
- Configure the Hold-up timer to 9999 (infinity) – this timer defines how long the RSMLT switch maintains forwarding for its peer

2.5.2 VRRP with ERS 8600

VRRP provides redundancy for the server's default gateway and can be utilized instead of RSMLT. It should be used for each configured VLAN that hosts end stations. Along with VRRP, Backup Master should be enabled on the Routing Switch Cluster to provide active-active routing and forwarding of traffic.




Figure 6 – ERS 8600 VRRP

When implementing VRRP, make sure to:

- Enable VRRP and Backup Master on each VLAN
- Configure VRRP priority higher than 100 (i.e. 110) to set VRRP Master and stagger VRRP Masters between ERS 8600's in the core
- Leave VRRP priority at default (100) for VRRP Backup
- Do not configure the virtual address as a physical interface that is used on any of the routing switches – use a third address, for example:
 - 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

2.6 Assumptions

This configuration uses a total of eight ERS 5000 series switches configured as stacks of four each, which are clustered together. Physical servers connect to this cluster. Two ERS 8600 switches are used for the routing function. RSMLT is proposed as the protocol of choice for creating redundant default gateways, although VRRP can also be used.

2.7 Objectives

This configuration provides a resilient, multi-switch connection for Hyper-V based servers. Fail-over and recovery times for NIC or link failures should be below 500 milliseconds. It should be possible that upgrades occur without any downtime.

2.8 Recommended Connectivity

Figure 7 depicts the topology used for testing and validation of the VMware solution. It is strongly recommended to distribute the NIC connections from the server to separate Ethernet switches in the stack for optimal resiliency when possible. A stack of four Ethernet switches is shown here, however, stacks less than four and greater than four (up to eight) are fully supported.




Figure 7 - Connectivity for VMware ESX Servers

2.9.2 Using RSMLT Layer 2 Edge

The recommended Layer 3 configuration using RSMLT Layer 2 Edge is show in Figure 9. The VMware servers are physically connected to the ERS 5000 Horizontal Stack Switch Cluster at Layer 2, however, the ERS 8600 core is performing all Layer 3 functions.




Figure 9 – Layer 3 Configuration Details using RSMLT for VMware Topology



Note that when RSMLT is used, there is no need to configure VRRP or any other routing layer redundancy protocol. RSMLT allows both L3 switches to be cognizant of each other's physical ip address which may have been configured as a next hop for some node prior to the Layer 3 switch. Both Layer 3 switches are able to discover if the other's interface is down and take appropriate measures to secure data path by responding to other's physical ip address until it comes back satisfactorily.

2.9.3 Using VRRP

Although RSMLT is Avaya's recommendation, VRRP can also be used for achieving gateway redundancy as shown in Figure 10. The VMware servers are physically connected to the ERS 5000 Horizontal Stack Switch Cluster at Layer 2, however, the ERS 8600 core is performing all Layer 3 functions.




Figure 10 – Layer 3 Configuration Details using VRRP for VMware Topology

2.9.4 Configure the Horizontal Stack Switch Cluster

2.9.4.1 Configure Management IP Address on ERS 5000 Horizontal Stack Switch Cluster

Configure Management IP Address on ERS 5000 stack 1 and on stack 2

```
5650TD(config)# ip default-gateway 10.160.192.254
5650TD(config)# ip address netmask 255.255.255.0
5650TD(config)# ip address stack 10.160.192.57

-----

5530-24TFD(config)# ip default-gateway 10.160.192.254
5530-24TFD(config)# ip address netmask 255.255.255.0
5530-24TFD(config)# ip address stack 10.160.192.54
```

2.9.4.2 Create VLANs on ERS 5000 Horizontal Stack Switch Cluster

Create VLANs on ERS 5000 stack 1 and on stack 2

```

5650TD(config)# vlan create 2 name "IST_VLAN" type port
5650TD(config)# vlan create 500 name "SQ_SMLT_TO_8600" type port
5650TD(config)# vlan create 510 name "SERV_OCS" type port
5650TD(config)# vlan create 550 name "HP_SERV_DNS" type port
5650TD(config)# vlan create 560 name "DELL_SERV_SQL" type port
5650TD(config)# vlan create 4000 name "VLAN #4000" type port
5650TD(config)# vlan ports 1/5,1/7,1/23,2/5,2/7,3/5,3/7,4/5,4/7,4/23 tagging tagAll
filter-untagged-frame-enable
5650TD(config)# vlan members remove 1 1/5,1/7,1/23,2/5,2/7,3/5,3/7,4/5,4/7,4/23
5650TD(config)# vlan members add 2 1/5,2/5,3/5,4/5
5650TD(config)# vlan members add 500 1/5,1/7,2/5,2/7,3/5,3/7,4/5,4/7
5650TD(config)# vlan members add 510 1/5,1/7,1/23,2/5,2/7,3/5,3/7,4/5,4/7,4/23
5650TD(config)# vlan members add 550 1/5,1/7,1/23,2/5,2/7,3/5,3/7,4/5,4/7,4/23
5650TD(config)# vlan members add 560 1/5,1/7,1/23,2/5,2/7,3/5,3/7,4/5,4/7,4/23
5650TD(config)# vlan members add 4000 1/1
5650TD(config)# vlan mgmt 4000

-----
5530-24TFD(config)# vlan create 2 name "IST_VLAN" type port
5530-24TFD(config)# vlan create 500 name "SQ_SMLT_TO_8600" type port
5530-24TFD(config)# vlan create 510 name "SERV_OCS" type port
5530-24TFD(config)# vlan create 550 name "HP_SERV_DNS" type port
5530-24TFD(config)# vlan create 560 name "DELL_SERV_SQL" type port
5530-24TFD(config)# vlan create 4000 name "VLAN #4000" type port
5530-24TFD(config)# vlan ports 1/5,1/7,1/23,2/5,2/7,3/5,3/7,4/5,4/7,4/23 tagging
tagAll filter-untagged-frames-enable
5530-24TFD(config)# vlan members remove 1 1/5,1/7,1/23,2/5,2/7,3/5,3/7,4/5,4/7, 4/23
5530-24TFD(config)# vlan members add 2 1/5,2/5,3/5,4/5
5530-24TFD(config)# vlan members add 500 1/5,1/7,2/5,2/7,3/5,3/7,4/5,4/7
5530-24TFD(config)# vlan members add 510 1/5,1/7,1/23,2/5,2/7,3/5,3/7,4/5,4/7, 4/23
5530-24TFD(config)# vlan members add 550 1/5,1/7,1/23,2/5,2/7,3/5,3/7,4/5,4/7, 4/23
5530-24TFD(config)# vlan members add 560 1/5,1/7,1/23,2/5,2/7,3/5,3/7,4/5,4/7, 4/23
5530-24TFD(config)# vlan members add 4000 1/1
5530-24TFD(config)# vlan mgmt 4000

```

2.9.4.3 Configure IST on ERS5000 Horizontal Switch Stack

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

```
5650TD(config)# ip routing
5650TD(config)# interface vlan 2
5650TD(config-if)# ip address 2.2.2.2 255.255.255.0 2
5650TD(config-if)# exit
-----
5530-24TFD(config)# ip routing
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 2.2.2.1 255.255.255.0 2
5530-24TFD(config-if)# exit
```

Create IST

```
5650TD(config)# interface mlt 1
5650TD(config-if)# ist peer-ip 2.2.2.1
5650TD(config-if)# ist vlan 2
5650TD(config-if)# ist enable
5650TD(config-if)# exit
-----
5530-24TFD(config)# interface mlt 1
5530-24TFD(config-if)# ist peer-ip 2.2.2.2
5530-24TFD(config-if)# ist vlan 2
5530-24TFD(config-if)# ist enable
5530-24TFD(config-if)# exit
```

2.9.4.4 Create MLTs on ERS 5000 Horizontal Stack Switch Cluster

Create MLTs on ERS 5000 stack 1 and on stack 2

```
5650TD(config)# mlt 1 name "IST"
5650TD(config)# mlt 1 member 1/5,2/5,3/5,4/5
5650TD(config)# mlt 1 learning disable
5650TD(config)# mlt 1 enable
5650TD(config)# mlt 2 name "SQ_SMLT_TO_8600"
5650TD(config)# mlt 2 member 1/7,2/7,3/7,4/7
5650TD(config)# mlt 2 learning disable
5650TD(config)# mlt 2 enable
-----
5530-24TFD(config)# mlt 1 name "IST"
5530-24TFD(config)# mlt 1 member 1/5,2/5,3/5,4/5
5530-24TFD(config)# mlt 1 learning disable
5530-24TFD(config)# mlt 1 enable
5530-24TFD(config)# mlt 2 name "SQ_SMLT_TO_8600"
5530-24TFD(config)# mlt 2 member 1/7,2/7,3/7,4/7
5530-24TFD(config)# mlt 2 learning disable
5530-24TFD(config)# mlt 2 enable
```

2.9.4.5 Configure VLACP on ERS 5000 switch clusters

Configure the VLACP MAC and enable VLACP globally on ERS5000 Stack 1 and Stack 2

```
5650TD(config)# vlacp macaddress 180.c200.f
5650TD(config)# vlacp enable
-----
5530-24TFD(config)# vlacp macaddress 180.c200.f
5530-24TFD(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*.



ERS 5510's do not support both Discard Untagged Frames and VLACP simultaneously.

Enable VLACP on the IST on ERS 5000 Stack 1 and Stack 2

```
5650TD(config)# interface FastEthernet all
5650TD(config-if)# vlacp port 1/5,2/5,3/5,4/5 timeout long
5650TD(config-if)# vlacp port 1/5,2/5,3/5,4/5 slow-periodic-time 10000
5650TD(config-if)# vlacp port 1/5,2/5,3/5,4/5 timeout-scale 3
5650TD(config-if)# vlacp port 1/5,2/5,3/5,4/5 enable
5650TD(config-if)# exit

-----

5530-24TFD(config)# interface FastEthernet all
5530-24TFD(config-if)# vlacp port 1/5,2/5,3/5,4/5 timeout long
5530-24TFD(config-if)# vlacp port 1/5,2/5,3/5,4/5 slow-periodic-time 10000
5530-24TFD(config-if)# vlacp port 1/5,2/5,3/5,4/5 timeout-scale 3
5530-24TFD(config-if)# vlacp port 1/5,2/5,3/5,4/5 enable
5530-24TFD(config-if)# exit
```

Enable VLACP on ports which will be used in SMLTs

```
5650TD(config)# interface FastEthernet all
5650TD(config-if)# vlacp port 1/7,2/7,3/7,4/7 timeout short
5650TD(config-if)# vlacp port 1/7,2/7,3/7,4/7 fast-periodic-time 500
5650TD(config-if)# vlacp port 1/7,2/7,3/7,4/7 timeout-scale 5
5650TD(config-if)# vlacp port 1/7,2/7,3/7,4/7 enable
5650TD(config-if)# exit

-----

5530-24TFD(config)# interface FastEthernet all
5530-24TFD(config-if)# vlacp port 1/7,2/7,3/7,4/7 timeout short
5530-24TFD(config-if)# vlacp port 1/7,2/7,3/7,4/7 fast-periodic-time 500
5530-24TFD(config-if)# vlacp port 1/7,2/7,3/7,4/7 timeout-scale 5
5530-24TFD(config-if)# vlacp port 1/7,2/7,3/7,4/7 enable
5530-24TFD(config-if)# exit
```



Both ends of the link must be configured for VLACP with the same parameters. The VLACP configuration for the ERS 8600 switches is shown later in this document.

2.9.4.6 Configure SMLT on ERS 5000 Horizontal Switch Stack Cluster

Configure SMLT on ERS 5000 (connection to ERS 8600 core)

```
5650TD(config)# interface mlt 4
5650TD(config-if)# smlt 4
5650TD(config-if)# exit

-----

5530-24TFD(config)# interface mlt 4
5530-24TFD(config-if)# smlt 4
5530-24TFD(config-if)# exit
```

2.9.4.7 Configure SLT on ERS 5000 Switch Cluster for Server Connections

Create SLT to Server

```
5650TD(config)#interface FastEthernet ALL
5650TD(config-if)#smlt port 1/23 33
5650TD(config-if)#smlt port 4/23 34
5650TD(config-if)#exit

-----

5530-24TFD(config)#interface FastEthernet ALL
5530-24TFD(config-if)#smlt port 1/23 33
5530-24TFD(config-if)#smlt port 4/23 34
5530-24TFD(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:

```
5650TD(config)# interface fastEthernet <port #>
5650TD(config-if)# spanning-tree learning disable
```

2.9.5 Configure ERS 8600 Switch Cluster Core (Layer 3)

2.9.5.1 Create VLANs on the ERS 8600 Switches

Create SMLT VLANs and IST VLAN on ERS8600-1 and ERS8600-2 Switches

```

C1-Dist-A:5# config vlan 500 create byport 1 name "SQ_SMLT_TO_55xx"
DC1-Dist-A:5# config vlan 500 fdb-entry aging-time 21601
DC1-Dist-A:5# config vlan 510 create byport 1 name "SERV_OCS"
DC1-Dist-A:5# config vlan 510 fdb-entry aging-time 21601
DC1-Dist-A:5# config vlan 550 create byport 1 name "HP_SERV_DNS"
DC1-Dist-A:5# config vlan 550 fdb-entry aging-time 21601
DC1-Dist-A:5# config vlan 560 create byport 1 name "DELL_SERV_SQL"
DC1-Dist-A:5# config vlan 560 fdb-entry aging-time 21601
DC1-Dist-A:5# config vlan 650 create byport 1 name "Load_Balancer"
DC1-Dist-A:5# config vlan 650 fdb-entry aging-time 21601
DC1-Dist-A:5# config vlan 2 create byport 1 "IST_VLAN"
-----
DC1-Dist-B:6# config vlan 500 create byport 1 name "SQ_SMLT_TO_55xx"
DC1-Dist-B:6# config vlan 500 fdb-entry aging-time 21601
DC1-Dist-B:5# config vlan 510 create byport 1 name "SERV_OCS"
DC1-Dist-B:5# config vlan 510 fdb-entry aging-time 21601
DC1-Dist-B:6# config vlan 550 create byport 1 name "HP_SERV_DNS"
DC1-Dist-B:6# config vlan 550 fdb-entry aging-time 21601
DC1-Dist-B:6# config vlan 560 create byport 1 name "DELL_SERV_SQL"
DC1-Dist-B:6# config vlan 560 fdb-entry aging-time 21601
DC1-Dist-B:6# config vlan 650 create byport 1 name "Load_Balancer"
DC1-Dist-B:6# config vlan 650 fdb-entry aging-time 21601
DC1-Dist-B:6# config vlan 2 create byport 1 "IST_VLAN"

```

Add IP addresses to the SMLT and IST VLANs

```

DC1-Dist-A:5# config vlan 2 ip create 10.2.12.1/255.255.255.252
DC1-Dist-A:5# config vlan 500 ip create 50.1.1.111/255.255.255.0
DC1-Dist-A:5# config vlan 510 ip create 51.0.1.111/255.255.255.0
DC1-Dist-A:5# config vlan 550 ip create 55.1.1.111/255.255.255.0
DC1-Dist-A:5# config vlan 560 ip create 56.1.1.111/255.255.255.252
DC1-Dist-A:5# config vlan 650 ip create 65.1.1.111/255.255.255.0

-----

DC1-Dist-B:6# config vlan 2 ip create 10.2.12.2/255.255.255.252
DC1-Dist-B:6# config vlan 500 ip create 50.1.1.112/255.255.255.0
DC1-Dist-B:5# config vlan 510 ip create 51.0.1.112/255.255.255.0
DC1-Dist-B:6# config vlan 550 ip create 55.1.1.112/255.255.255.0
DC1-Dist-B:6# config vlan 560 ip create 56.1.1.112/255.255.255.252
DC1-Dist-B:6# config vlan 650 ip create 65.1.1.112/255.255.255.0

```

2.9.5.2 Create ERS 8600 Switch Cluster

Create MLT 1 for IST

```

DC1-Dist-A:5# config mlt 1 create
DC1-Dist-A:5# config mlt 1 name "IST_MLT"
DC1-Dist-A:5# config mlt 1 add ports 2/1,2/2
DC1-Dist-A:5# config mlt 1 perform-tagging enable
DC1-Dist-A:5# config vlan 2 add-mlt 1

-----

DC1-Dist-B:6# config mlt 1 create
DC1-Dist-B:6# config mlt 1 name "IST_VLAN"
DC1-Dist-B:6# config mlt 1 add ports 2/1,2/2
DC1-Dist-B:6# config mlt 1 perform-tagging enable
DC1-Dist-B:6# config vlan 2 add-mlt 1

```

Create IST

```
DC1-Dist-A:5# config mlt 1 ist create ip 10.2.12.2 vlan-id 2
DC1-Dist-A:5# config mlt 1 ist enable

-----

DC1-Dist-B:6# config mlt 1 ist create ip 10.2.12.1 vlan-id 2
DC1-Dist-B:6# config mlt 1 ist enable
```

Create SMLT on ERS 8600 Distribution Switches

```
DC1-Dist-A:5# config mlt 21 create
DC1-Dist-A:5# config mlt 21 add ports 3/1-3/2,3/7-3/8
DC1-Dist-A:5# config mlt 21 name "SQ_SMLT_TO_56xx"
DC1-Dist-A:5# config mlt 21 perform-tagging enable
DC1-Dist-A:5# config mlt 21 smlt create smlt-id 21
DC1-Dist-A:5# config mlt 24 create
DC1-Dist-A:5# config mlt 24 add ports 3/17-3/18
DC1-Dist-A:5# config mlt 24 name "TRI_SMLT_LB"
DC1-Dist-A:5# config mlt 24 perform-tagging enable
DC1-Dist-A:5# config mlt 24 smlt create smlt-id 24

-----

DC1-Dist-B:6# config mlt 21 create
DC1-Dist-B:6# config mlt 21 add ports 3/1-3/2,3/7-3/8
DC1-Dist-B:6# config mlt 21 name "SQ_SMLT_TO_56xx"
DC1-Dist-B:6# config mlt 21 perform-tagging enable
DC1-Dist-B:6# config mlt 21 smlt create smlt-id 21
DC1-Dist-B:6# config mlt 24 create
DC1-Dist-B:6# config mlt 24 add ports 3/17-3/18
DC1-Dist-B:6# config mlt 24 name "TRI_SMLT_LB"
DC1-Dist-B:6# config mlt 24 perform-tagging enable
DC1-Dist-B:6# config mlt 24 smlt create smlt-id 24
```

Add VLANs to IST and SMLT

```
DC1-Dist-A:5# config vlan 500 add-mlt 1
DC1-Dist-A:5# config vlan 500 add-mlt 21
DC1-Dist-A:5# config vlan 510 add-mlt 1
DC1-Dist-A:5# config vlan 510 add-mlt 21
DC1-Dist-A:5# config vlan 550 add-mlt 1
DC1-Dist-A:5# config vlan 550 add-mlt 21
DC1-Dist-A:5# config vlan 560 add-mlt 1
DC1-Dist-A:5# config vlan 560 add-mlt 21
DC1-Dist-A:5# config vlan 650 add-mlt 1
DC1-Dist-A:5# config vlan 650 add-mlt 24
```

```
-----
DC1-Dist-B:6# config vlan 500 add-mlt 1
DC1-Dist-B:6# config vlan 500 add-mlt 21
DC1-Dist-B:6# config vlan 510 add-mlt 1
DC1-Dist-B:6# config vlan 510 add-mlt 21
DC1-Dist-B:6# config vlan 550 add-mlt 1
DC1-Dist-B:6# config vlan 550 add-mlt 21
DC1-Dist-B:6# config vlan 560 add-mlt 1
DC1-Dist-B:6# config vlan 560 add-mlt 21
DC1-Dist-B:6# config vlan 650 add-mlt 1
DC1-Dist-B:6# config vlan 650 add-mlt 24
```

2.9.5.3 Enable VLACP on the ERS 8600 Switches

Enable VLACP globally

```
DC1-Dist-A:5# config vlacp enable
```

```
DC1-Dist-B:6# config vlacp enable
```

Enable VLACP on IST

```
DC1-Dist-A:5# config ethernet 2/1,2/2 vlacp macaddress 01:80:c2:00:00:0f
```

```
DC1-Dist-A:5# config ethernet 2/1,2/2 vlacp slow-periodic-time 10000
```

```
DC1-Dist-A:5# config ethernet 2/1,2/2 vlacp timeout long
```

```
DC1-Dist-A:5# config ethernet 2/1,2/2 vlacp timeout-scale 3
```

```
DC1-Dist-A:5# config ethernet 2/1,2/2 vlacp enable
```

```
DC1-Dist-B:6# config ethernet 2/1,2/2 vlacp macaddress 01:80:c2:00:00:0f
```

```
DC1-Dist-B:6# config ethernet 2/1,2/2 vlacp slow-periodic-time 10000
```

```
DC1-Dist-B:6# config ethernet 2/1,2/2 vlacp timeout long
```

```
DC1-Dist-B:6# config ethernet 2/1,2/2 vlacp timeout-scale 3
```

```
DC1-Dist-B:6# config ethernet 2/1,2/2 vlacp enable
```

Enable VLACP on SMLT

```
DC1-Dist-A:5# config ethernet 3/1,3/2,3/7,3/8,3/17,3/18 vlacp macaddress 01:80:c2:00:00:0f
```

```
DC1-Dist-A:5# config ethernet 3/1,3/2,3/7,3/8,3/17,3/18 vlacp fast-periodic-time 500
```

```
DC1-Dist-A:5# config ethernet 3/1,3/2,3/7,3/8,3/17,3/18 vlacp timeout short
```

```
DC1-Dist-A:5# config ethernet 3/1,3/2,3/7,3/8,3/17,3/18 vlacp timeout-scale 5
```

```
DC1-Dist-A:5# config ethernet 3/1,3/2,3/7,3/8,3/17,3/18 vlacp enable
```

```
DC1-Dist-B:6# config ethernet 3/1,3/2,3/7,3/8,3/17,3/18 vlacp macaddress 01:80:c2:00:00:0f
```

```
DC1-Dist-B:6# config ethernet 3/1,3/2,3/7,3/8,3/17,3/18 vlacp fast-periodic-time 500
```

```
DC1-Dist-B:6# config ethernet 3/1,3/2,3/7,3/8,3/17,3/18 vlacp timeout short
```

```
DC1-Dist-B:6# config ethernet 3/1,3/2,3/7,3/8,3/17,3/18 vlacp timeout-scale 5
```

```
DC1-Dist-B:6# config ethernet 3/1,3/2,3/7,3/8,3/17,3/18 vlacp enable
```

2.9.5.4 Configure Layer 3 on the ERS 8600 Switches

2.9.5.4.1 Configuration for RSMLT

VLAN Name	VLAN ID	IP Address	Hold Down Timer	Hold Up Timer
<i>SQ_SMLT_TO_56xx</i>	500	50.1.1.111 50.1.1.112	60	180
<i>SERV_OCS</i>	510	51.0.1.111 51.0.1.112	60	180
<i>HP_SERV_DNS</i>	550	55.1.1.111 55.1.1.112	60	180
<i>DELL_SERV_SQL</i>	560	56.1.1.111 56.1.1.112	60	180
<i>TRI_SMLT_LB</i>	650	65.1.1.111 65.1.1.112	60	180

Enable RSMLT

```

DC1-Dist-A:5# config vlan 500 ip rsmlt enable
DC1-Dist-A:5# config vlan 500 ip rsmlt holdup-timer 9999
DC1-Dist-A:5# config vlan 510 ip rsmlt enable
DC1-Dist-A:5# config vlan 510 ip rsmlt holdup-timer 9999
DC1-Dist-A:5# config vlan 550 ip rsmlt enable
DC1-Dist-A:5# config vlan 550 ip rsmlt holdup-timer 9999
DC1-Dist-A:5# config vlan 560 ip rsmlt enable
DC1-Dist-A:5# config vlan 560 ip rsmlt holdup-timer 9999
DC1-Dist-A:5# config vlan 650 ip rsmlt enable
DC1-Dist-A:5# config vlan 650 ip rsmlt holdup-timer 9999
DC1-Dist-A:5# config ip rsmlt rsmlt-edge-support enable
-----
DC1-Dist-B:6# config vlan 500 ip rsmlt enable
DC1-Dist-B:6# config vlan 500 ip rsmlt holdup-timer 9999
DC1-Dist-B:6# config vlan 510 ip rsmlt enable
DC1-Dist-B:6# config vlan 510 ip rsmlt holdup-timer 9999
DC1-Dist-B:6# config vlan 550 ip rsmlt enable
DC1-Dist-B:6# config vlan 550 ip rsmlt holdup-timer 9999
DC1-Dist-B:6# config vlan 560 ip rsmlt enable
DC1-Dist-B:6# config vlan 560 ip rsmlt holdup-timer 9999
DC1-Dist-B:6# config vlan 650 ip rsmlt enable
DC1-Dist-B:6# config vlan 650 ip rsmlt holdup-timer 9999
DC1-Dist-B:6# config ip rsmlt rsmlt-edge-support enable

```

2.9.5.4.2 Configuration for VRRP

VRRP can also be used for providing default gateway redundancy, however RSMLT configuration shown above is our recommendation. We show the VRRP configuration if it is preferred.

VLAN Name	VLAN ID	VRRP IP Address	Dist-A Priority	Dist-B Priority	Adv Interval	Hold down Timer
<i>SQ_SMLT_TO_56xx</i>	500	50.1.1.113	200	Default	10	60
<i>SERV_OCS</i>	510	51.0.1.113	Default	200	10	60
<i>HP_SERV_DNS</i>	550	55.1.1.113	Default	200	10	60
<i>DELL_SERV_SQL</i>	560	56.1.1.113	200	Default	10	60
<i>TRI_SMLT_LB</i>	650	65.1.1.113	Default	200	10	60

Enable VRRP and Backup Master

```

DC1-Dist-A:5# config vlan 500 ip vrrp 50 address 50.1.1.113
DC1-Dist-A:5# config vlan 500 ip vrrp 50 backup-master enable
DC1-Dist-A:5# config vlan 500 ip vrrp 50 priority 200
DC1-Dist-A:5# config vlan 500 ip vrrp 50 adver-int 10
DC1-Dist-A:5# config vlan 500 ip vrrp 50 holddown-timer 60
DC1-Dist-A:5# config vlan 500 ip vrrp 50 enable
DC1-Dist-A:5# config vlan 510 ip vrrp 51 address 51.0.1.113
DC1-Dist-A:5# config vlan 510 ip vrrp 51 backup-master enable
DC1-Dist-A:5# config vlan 510 ip vrrp 51 adver-int 10
DC1-Dist-A:5# config vlan 510 ip vrrp 51 holddown-timer 60
DC1-Dist-A:5# config vlan 510 ip vrrp 51 enable
DC1-Dist-A:5# config vlan 550 ip vrrp 55 address 55.1.1.113
DC1-Dist-A:5# config vlan 550 ip vrrp 55 backup-master enable
DC1-Dist-A:5# config vlan 550 ip vrrp 55 adver-int 10
DC1-Dist-A:5# config vlan 550 ip vrrp 55 holddown-timer 60
DC1-Dist-A:5# config vlan 550 ip vrrp 55 enable
DC1-Dist-A:5# config vlan 560 ip vrrp 56 address 56.1.1.113
DC1-Dist-A:5# config vlan 560 ip vrrp 56 backup-master enable
DC1-Dist-A:5# config vlan 560 ip vrrp 56 priority 200
DC1-Dist-A:5# config vlan 560 ip vrrp 56 adver-int 10

```

```

DC1-Dist-A:5# config vlan 560 ip vrrp 56 holddown-timer 60
DC1-Dist-A:5# config vlan 560 ip vrrp 56 enable
DC1-Dist-A:5# config vlan 650 ip vrrp 65 address 65.1.1.113
DC1-Dist-A:5# config vlan 650 ip vrrp 65 backup-master enable
DC1-Dist-A:5# config vlan 650 ip vrrp 65 adver-int 10
DC1-Dist-A:5# config vlan 650 ip vrrp 65 holddown-timer 60
DC1-Dist-A:5# config vlan 650 ip vrrp 65 enable
-----
DC1-Dist-B:6# config vlan 500 ip vrrp 50 address 50.1.1.113
DC1-Dist-B:6# config vlan 500 ip vrrp 50 backup-master enable
DC1-Dist-B:6# config vlan 500 ip vrrp 50 adver-int 10
DC1-Dist-B:6# config vlan 500 ip vrrp 50 holddown-timer 60
DC1-Dist-B:6# config vlan 500 ip vrrp 50 enable
DC1-Dist-B:5# config vlan 510 ip vrrp 51 address 51.0.1.113
DC1-Dist-B:5# config vlan 510 ip vrrp 51 backup-master enable
DC1-Dist-B:6# config vlan 510 ip vrrp 51 priority 200
DC1-Dist-B:5# config vlan 510 ip vrrp 51 adver-int 10
DC1-Dist-B:5# config vlan 510 ip vrrp 51 holddown-timer 60
DC1-Dist-B:5# config vlan 510 ip vrrp 51 enable
DC1-Dist-B:6# config vlan 550 ip vrrp 55 address 55.1.1.113
DC1-Dist-B:6# config vlan 550 ip vrrp 55 backup-master enable
DC1-Dist-B:6# config vlan 550 ip vrrp 55 priority 200
DC1-Dist-B:6# config vlan 550 ip vrrp 55 adver-int 10
DC1-Dist-B:6# config vlan 550 ip vrrp 55 holddown-timer 60
DC1-Dist-B:6# config vlan 550 ip vrrp 55 enable
DC1-Dist-B:6# config vlan 560 ip vrrp 56 address 56.1.1.113
DC1-Dist-B:6# config vlan 560 ip vrrp 56 backup-master enable
DC1-Dist-B:6# config vlan 560 ip vrrp 56 adver-int 10
DC1-Dist-B:6# config vlan 560 ip vrrp 56 holddown-timer 60
DC1-Dist-B:6# config vlan 560 ip vrrp 56 enable
DC1-Dist-B:6# config vlan 650 ip vrrp 65 address 65.1.1.113
DC1-Dist-B:6# config vlan 650 ip vrrp 65 backup-master enable
DC1-Dist-B:6# config vlan 650 ip vrrp 65 priority 200
DC1-Dist-B:6# config vlan 650 ip vrrp 65 adver-int 10
DC1-Dist-B:6# config vlan 650 ip vrrp 65 holddown-timer 60
DC1-Dist-B:6# config vlan 650 ip vrrp 65 enable

```

2.9.6 Create and Configure VMs using vSphere Client


In this section, we detail the steps needed for creating and configuring virtual machines using the vSphere client. We describe the method for creating a 'Test' virtual machine. This method can be used repeatedly to create various virtual machines used in the example scenario.

Overview of steps


1. Create Virtual Machine
2. Create Virtual Switch for NIC Teaming
3. Add adapter ports to the virtual switch
4. Add virtual switch to the VM configuration
5. Power on the VM
6. Configure networking and other applications in VM

2.9.6.1 Creation of Virtual Machine


Start VMware vSphere Client from the Desktop icon or **Start->All Programs->VMware->VMware vSphere Client**.




To create a new Virtual Machine, right click on the server **convergence-esx1.us.nortel.com** and select **New Virtual Machine** and the **Create New Virtual Machine** window will come up.




In the **Create New Virtual Machine Window**, it gives an option to select **Typical** or **Custom Configuration**. Click on the **Next** button after selecting the **Configuration Type** to continue the creation of a new Virtual Machine.




Specify the **Name** and **Location** for the Virtual Machine that is being created. Click the **Next** button to proceed.




Select a **Datastore** for the Virtual Machine from the list of the datastores that are configured for the server. Click the **Next** button to proceed.




Specify the **Guest Operating System** that has to be used in the Virtual Machine from the list of options that is available. Click on **Next** button to proceed.



Specify the **Virtual Hard Disk** size for the Virtual machine (the default size is 8 GB). Click on **Next** button to proceed.




The next screen gives the summary of the settings for the new Virtual machine. Click on the **Finish** button to complete the creation.




2.9.6.2 Creation of Virtual Switch for NIC Teaming


In the vSphere Client window for the convergence-esx1.us.nortel.com server, select the **Configuration** tab and click on the **Networking** option. It will show the **Virtual Switch** configuration.




For adding a new Virtual Switch, click on **Add Networking** at the top right corner to open the **Add Network Wizard**. The **Add Network Wizard** first displays the type of **Connection Type** (Virtual Machine/VMkernel/Service Console) to be used. Select Virtual Machine and click on the **Next** button at the bottom of the wizard.




Select the **Network Access** to be used. Choose the **Create a virtual switch** option to create a new one. Click on the **Next** button at the bottom of the window to proceed.



Specify a name for the **Virtual Machine Port** group and add a **VLAN Id** if tagging needs to be done. Click on the **Next** button.




Click the **Finish** button to complete this part of the **Virtual Switch** creation.




2.9.6.3 Addition of Adaptor Ports to the Virtual Switch


Under vSphere Client, select the **Networking** option under the **Configuration** tab of the convergence-esx1.us.nortel.com server. Select **Properties** for the **Virtual Switch** that was created in the above step.




This screen will show the **Port** group that was created and its properties: Number of ports, Security, Traffic Shaping, Failover and Load Balancing configuration.




The properties of the **Virtual Switch** can be changed by clicking the **Edit** button in the window.




Click on the **Network Adapters** tab to add/edit network adapter information



Click on **Add** to add a new adapter to the group. Select required ports from the options available.




Once added, the **Edit** button can be used to make changes to port configuration. Click **Close** to complete this step. The **Virtual Switch** is now configured.




2.9.6.4 Adding Virtual Switch to Virtual Manager


Under the ESX (convergence-es1.us.nortel.com) server – select the new **VM** that was added and needs configuration.




Go the **Summary** tab as shown below



Click on **Edit Settings** and select **Network Adapter 1** option




Under **Network Connection**, select the **Virtual Switch** port group that was created




2.9.6.5 Powering on the VM

The **VM** can be now be powered on by right clicking it as shown below. This will cause the selected OS to be installed for the VM.




To view the **Summary** of the **Configuration** on the VM, click on the **Summary** tab.



2.9.6.6 Configure Networking and other Applications in VM

The network and other configurations on the VM can be done by selecting the **Console** tab or by right clicking the VM and selecting the **Open Console** option. This can be used to set the IP address, default gateway and DNS IP address for the VM. It can also be used to install applications such as OCS.



2.10 Verification

To test this configuration:

- From each client ensure that you can ping the load balancer virtual ip address.
- From each VM, ensure that you can ping the gateway on the 8600.
- From each VM, ensure that you can ping each other VM.
- From the client connect to the OCS server and ensure that the connectivity is established with low latency.
- While a connection between the client and the server is up, pull out links one by one until a path containing at least one link on each hop remains. Ensure that the session established between the client and the server was not disrupted.
- While a connection between the client and the server is up, re-connect all of the disconnected links one by one. Ensure that the session established between the client and the server was not disrupted.

3. 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.

3.1 Getting technical documentation

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

3.2 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.

3.3 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.

3.4 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.