Virtual Services Platform 9000
and
Virtual Services Platform 7000

Engineering

>Avaya Virtual Services Platform 9000 and Avaya Virtual Services Platform 7000 with Coraid EtherDrive SRX-Series Storage Appliances Technical Configuration Guide

Avaya Networking

Document Date: August 2011

Document Number: NN48500-629

Document Version: 1.0

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

This Technical Configuration Guide describes a solution comprised of the Avaya Virtual Services Platform 9000, the Avaya Virtual Services Platform 7000, and Coraid EtherDrive storage appliances.

Information in this Technical Configuration Guide has been obtained through Avaya Networking interoperability testing and additional technical discussions. Testing was conducted at the Avaya Networking Test Lab.

Acronym Key

Throughout this guide the following acronyms will be used:

AoE	ATA over Ethernet
ATA	Advanced Technology Attachment
HBA	Host Bus Adapter
LACP	Link Aggregation Control Protocol
LUN	Logical Unit Number
MLT	Multi-Link Trunking
LUN	Logical Unit Number
RAID	Redundant Array of Independent Disks
SAN	Storage Area Network
SAS	Serial Attached SCSI
SATA	Serial ATA
SCSI	Small Computer System Interface
SSD	Solid State Disk Drive

Table of Contents

Figures

Tables

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

Bold text indicates emphasis.

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

```
VSP_CoreA:1#show running-config
```

Output examples from Avaya devices are displayed in a Lucida Console font:

```
VSP_CoreA:1#show sys-info
General Info :

        SysDescr     : VSP-9012 (3.0.2.0 GA)
        SysName      : VSP_CoreA
        SysUpTime    : 28 day(s), 17:13:23
        SysContact   : http://support.avaya.com/
        SysLocation  : 211 Mt. Airy Road,Basking Ridge,NJ 07920

Chassis Info:

        Chassis            : 9012
        Serial#            :
        H/W Revision       :
        H/W Config         :
        NumSlots           : 12
        NumPorts           : 121
        BaseMacAddr        : 00:24:7f:9e:a0:00
        MacAddrCapacity    : 4096
        MgmtMacAddr        : 00:24:7f:9e:a3:fd
        System MTU         : 1950
```

1. Introduction

This Technical Configuration Guide describes an Ethernet storage solution comprised of the following components:

- Avaya Virtual Services Platform 9000
- Avaya Virtual Services Platform 7000
- Coraid EtherDrive storage appliances.

During interoperability testing, two Coraid SRX EtherDrive SAN storage appliances successfully received and stored data transmitted from Avaya Virtual Services Platform 7000 aggregation switches through the Avaya Virtual Services Platform 9000 core switches over 10 GbE links.

1.1 Avaya Virtual Services Platform 9000

The Avaya Virtual Services Platform 9000 (VSP 9000) is a new Ethernet Switching platform for Enterprise Campus environments and Enterprise Data Centers. This platform offers an unmatched switching architecture that scales from an initial 8.4 Terabits per second to an industry-leading 27 Terabits per second. The VSP 9000 delivers substantial performance and scalability, with immediate support for very high-density 1 and 10 GbE, in addition to being future-ready for the emerging 40 and 100 GbE standards. The fully scalable architecture helps ensure that network capacity seamlessly scales in line with performance requirements, without complex or expensive re-engineering.

The Avaya VSP 9000 architecture is ultra-reliable and has the following features that help ensure uninterrupted business operations:

- Fully redundant hardware, including the control processor and switch fabric modules, that ensure no single point-of-failure
- Switch Clustering delivering deterministic millisecond failover resiliency for instantaneous recovery from any individual failure or during maintenance without impacting user applications
- Layer 2 and Layer 3 network virtualization services providing support for multiple customers and user groups on the same platform
- Network failover in less than 20 milliseconds with instantaneous re-route across all ports to minimize packet loss
- “In-service control plane integrity check" and “rapid failure detection and recovery of data path” for system-level health check and self-healing capabilities
- Hitless patching eliminating the requirement to reload the complete system image, thereby minimizing maintenance down time
- Flight Recorder style logging capability to help with continuous real-time monitoring of internal control message flows
- Key Health Indicators to provide system operators with a view of system health on all levels: OS, system applications /protocols I/O modules, ports and the forwarding path
- Ability to remotely update flash images
- Avaya Virtual Services support using IEEE Shortest Path Bridging – de-couples physical infrastructure from logical provisioning and ensures predictability for all network services

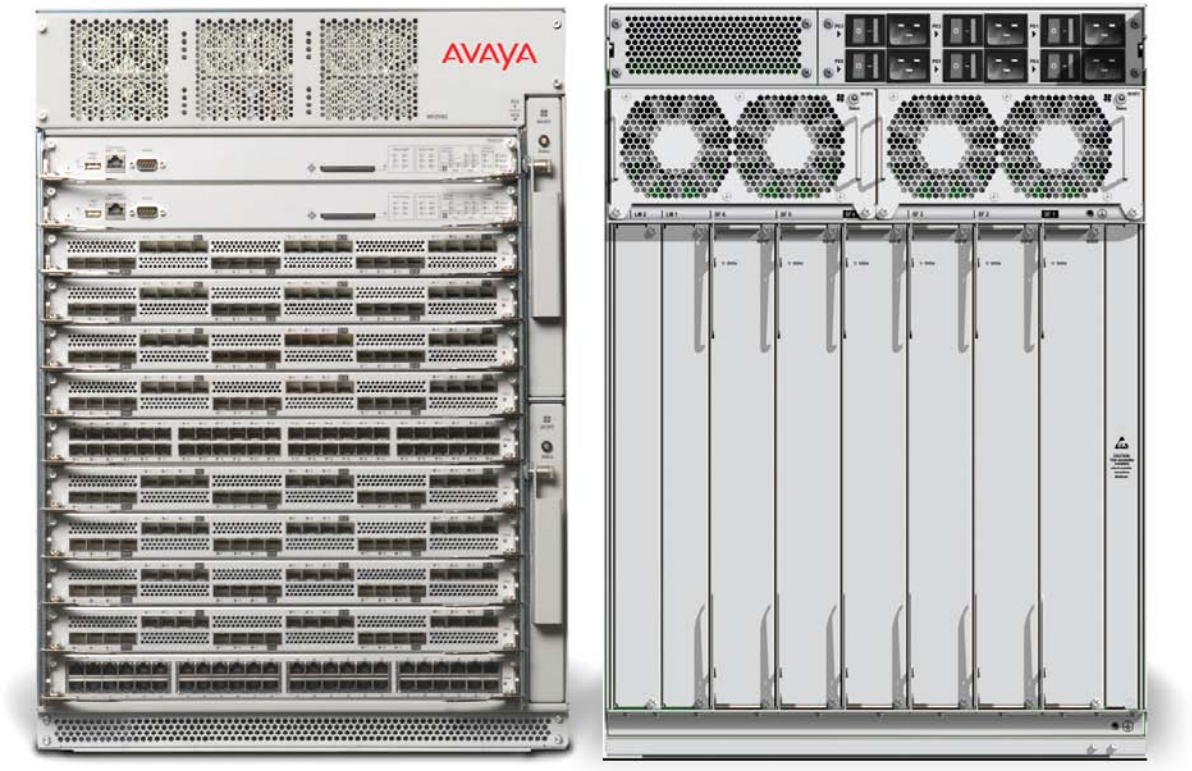

Figure 1 – Avaya VSP 9000 Ethernet Switch

1.2 Avaya Virtual Services Platform 7000

The Avaya Virtual Services Platform (VSP 7000) is a new family of 1/10Gigabit, Top of Rack, Ethernet Switches. These high-density, high-capacity switches provide a high performance forwarding engine for data centers aggregation and small to medium core switches. The following is a list of some of the Avaya VSP 7000 features:

- 1RU stackable switch with class-leading switching performance of over 1.2Tbps
- Data center grade hardware that supports front-to-back or back-to-front cooling
- 5th generation ASIC technology for future proof feature requirements
- 24 ports of SFP+ supporting both/either 1 and 10 GbE
- Media Dependant Adaptor (MDA) for a range of high-speed expansion options
- SFP+ connectivity to connect at 1 Gigabit or 10 Gigabit speeds
- Future-ready with flexible support for 40Gbps, 100Gbps Ethernet and Fibre Channel
- Support network-wide fabric-based Virtualized Services and Lossless environments
- Dual, hot-swappable AC or DC power supplies and fan trays for always-on high-performance

The Avaya VSP 7000 switch is designed for Enterprise customers requiring high density, high performance 10 Gigabit connectivity. In a high-performance Data Center, the Avaya VSP 7000 can serve as a Top-of-Rack Switch. In a network with an existing Core Switch deployment, it can provide a cost-effective 10 Gigabit Ethernet fan-out capability. In a Campus distribution layer, it can deliver flexible connectivity and consolidation options.

Figure 2 – Avaya VSP 7000 Ethernet Switch

1.3 Coraid EtherDrive SRX Series

Coraid is redefining the fundamental economics of storage with Ethernet SAN solutions that provide enterprises of all sizes with a flexible tier of high-performance, scale-out storage. The Coraid EtherDrive SRX Series storage arrays deliver performance up to 1,800 MB/sec, with multiple 10 GbE connections per shelf, providing a 5-8x price performance advantage over other SAN storage solutions. EtherDrive SRX scales to multiple petabytes, uniquely delivering the fundamental building blocks necessary to build the next generation SAN infrastructure for a range of applications including server virtualization, high-performance computing, and cloud storage.

Coraid uses RAID technology as a method of logically treating several hard drives as one unit to improve performance and/or provide redundancy. Coraid EtherDrive SRX-Series storage appliances support all standard RAID types (0, 1, 5, 6, and 10). The Coraid EtherDrive SRX-Series appliances are block storage RAID devices with front loading, hot-swappable SATA, SAS, and SSD disk drives. Each EtherDrive SRX appliance can be used individually or in RAID sets.

Figure 3 – Coraid EtherDrive SRX Storage Array

1.4 Interoperability and Resiliency Testing

The Avaya / Coraid solution testing focused primarily on verifying interoperability and resiliency simulating various fault conditions to ensure high availability. During interoperability testing, the Avaya VSP 7000s were used as Switch Clustering edge devices in conjunction with two Avaya VSP 9000 core devices and communicated with two Coraid EtherDrive SRX-Series storage appliances. This verified the resiliency and fast failover of the Switch Clustering solution of the Avaya VSP 9000.

The Coraid EtherDrive SRX-Series uses the ATA over Ethernet (AoE) protocol to communicate between the logical storage devices and servers across a standard Ethernet network. AoE enables you to share disk drives through a standard Ethernet network. AoE arranges the communication that would normally take place between a server and a disk drive into data packets called datagrams and sends these across Layer 2 Ethernet with minimal overhead. Datagrams are addressed to storage devices using their Ethernet MAC addresses.

AoE does not run over high level networking protocols like IP so AoE datagrams cannot be routed. The data packets can travel across the switches that make up an Ethernet LAN, but routers cannot send them to another network and devices outside of the LAN cannot communicate with them. This provides an inherent layer of security. AoE is designed to run on a datagram networking protocol such as Ethernet, which makes a best effort attempt to deliver datagrams. AoE has developed a sophisticated congestion avoidance algorithm to maximize throughput while avoiding dropped frames. AoE can quickly recover from lost datagrams on the network due to congestion so it is guaranteed to not lose data.

For technical support, access to documentation, frequently asked questions, AoE tools for Linux, and contact telephone numbers on Coraid EtherDrive SRX-Series storage appliances, contact Coraid's technical support at their Web site: http://www.coraid.com/support/customer_support.

2. Reference Configuration

Figure 4 shows a sample storage area network (SAN) configuration with two Avaya VSP 9000s, two Avaya VSP 7000s, five PC servers, and two Coraid EtherDrive SRX-Series storage appliances. All communication between these devices use 10 GbE links in a Layer 2 topology. All of the 10 GbE connections use SFP+ connectors with multimode fiber.

Coraid recommends that you isolate the SAN from other network traffic.

To verify data integrity of the Coraid EtherDrive SRX-Series storage appliances, the Avaya Networking Test Lab used the Open Source software, Vdbench, which is a disk and tape I/O workload generator for directly attached and network connected storage devices.

For Vdbench downloads and documentation, go to http://sourceforge.net/projects/vdbench/.

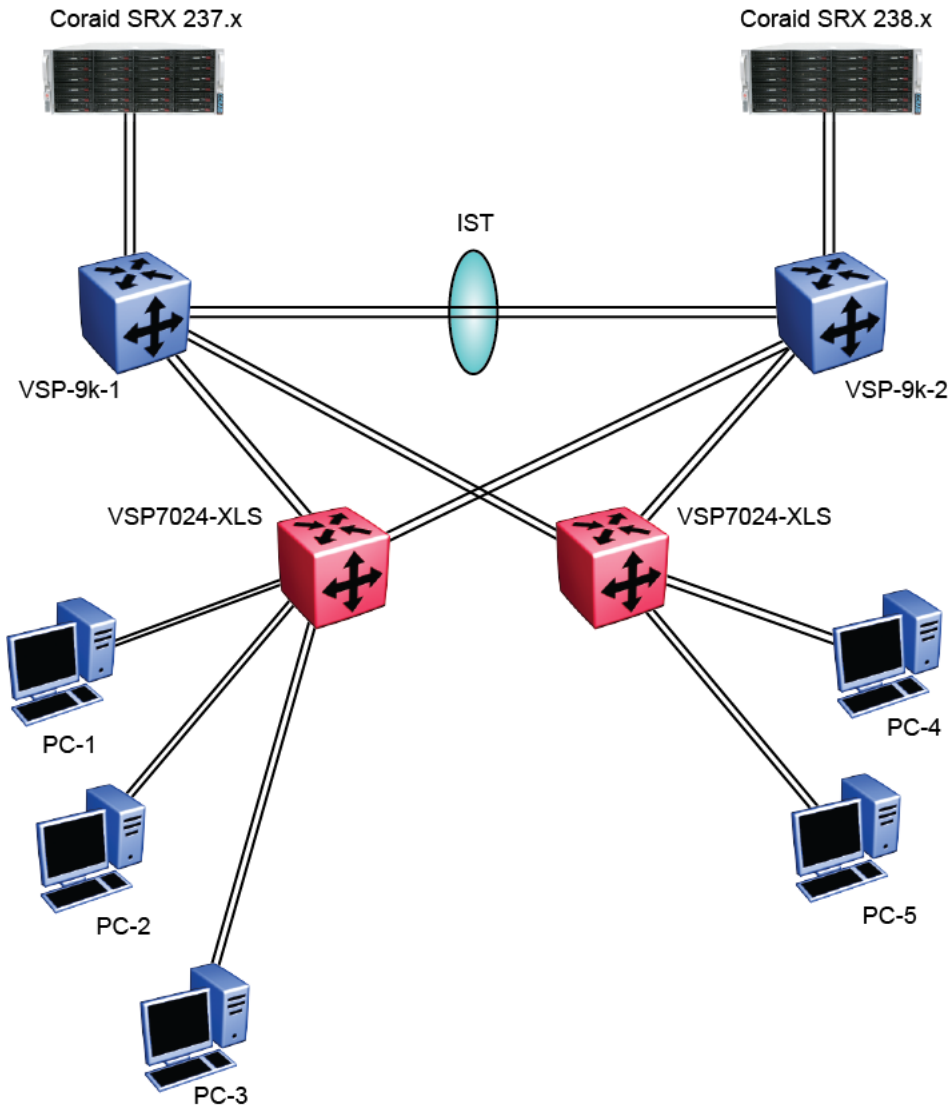

Figure 4 – SAN Test Bed Configuration

3. Equipment and Software Validated

The following equipment and software were used for the sample configuration provided:

Equipment	Software
• 2 Avaya VSP 9000 Ethernet Switches	Software Release GA version 3.0.2
• 2 Avaya VSP7024-XLS Ethernet Switches	Software Release 10.0.0b36
• 5 PC 1U servers with Coraid Host Bus Adapter (HBA) network cards installed	Coraid HBA build 48 CentOS 5.5 64-bit
• 2 SRX-Series storage appliances each containing: o 2 dual-port 10GbE network interface cards o 25 300GB SATA drives	CorOS 5.1.4
• 5 Coraid 10GbE, dual port, SFP+ HBA Cards (1 per server)	
	Vdbench verification software

Table 1 – Hardware and Software Versions

4. Storage Area Network Configuration

This section provides configuration details for the following SAN components:

- Avaya VSP 9000
- Avaya VSP 7000
- PC servers
- Coraid EtherDrive SRX-Series storage appliances

4.1 Avaya Virtual Services Platform 9000 Configuration

This section describes how to configure the Avaya VSP 9000. To allow for full utilization of the network for SAN storage, Coraid recommends isolating SAN storage networks from other network traffic.

Note – In this test setup, STP is disabled on all connected ports for faster transition time when a port is unplugged and plugged during. Also, the specific VLANs and ports used have no special significance and therefore can be changed at the user's discretion.

4.1.1 Configuring VLANs using the ACLI

Configure VLAN 100 and ensure that all ports connected to the Coraid HBAs and SRX-Series appliances are in VLAN 100. The following steps show how to configure the VLAN using the ACLI:

1 Enter the *Privilege Exec* context:

```
VSP9000:1>enable
```

2 Enter the *Global Configuration* context:

```
VSP9000:1# configure terminal
```

3 Create the port based VLAN 100:

```
VSP9000:1(config)# vlan create 100 name SAN type port-mstprstp 1
```

4 Remove the ports from the default VLAN 1:

```
VSP9000:1(config)# vlan members remove 1 3/1-3/9,6/1-6/9
```

5 Add ports to the VLAN 100:

```
VSP9000:1(config)# vlan members add 100 3/1-3/9,6/1-6/9
```

4.1.2Configuring VLANs using Enterprise Device Manager

Configure VLAN 100 and ensure that all ports connected to the Coraid HBAs and SRX-Series appliances are in VLAN 100. The following steps show how to configure the VLAN using Enterprise Device Manager (EDM):

1 Select *Configuration > VLAN > VLANs > Basic*. Double click on the *PortMembers* for *VLAN 1*:

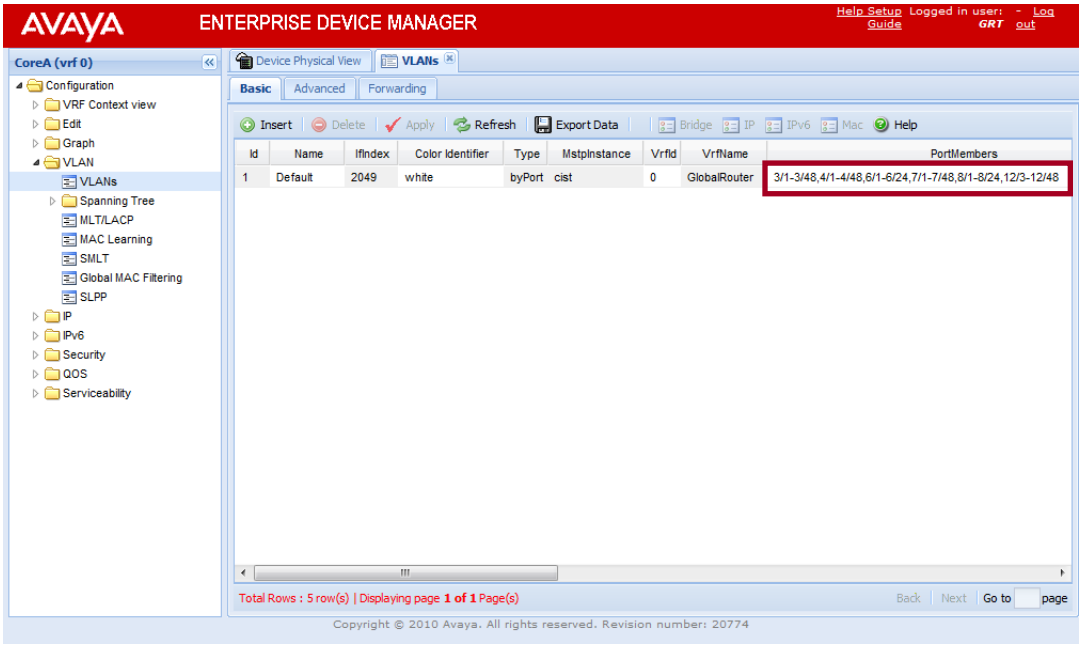

2 Uncheck ports *1/1-1/9* and *6/1-6/9* then click *OK*:

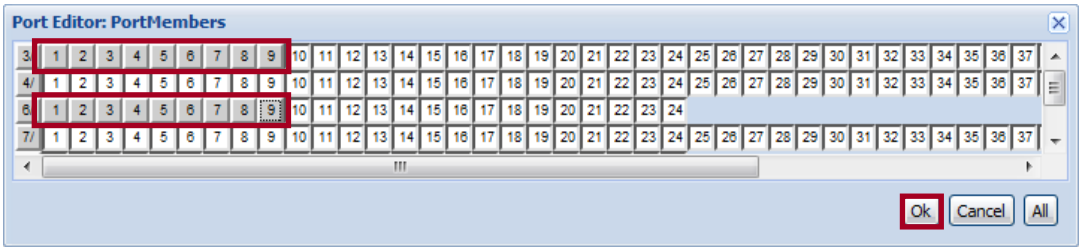

3 Create a new VLAN by clicking *Insert*:

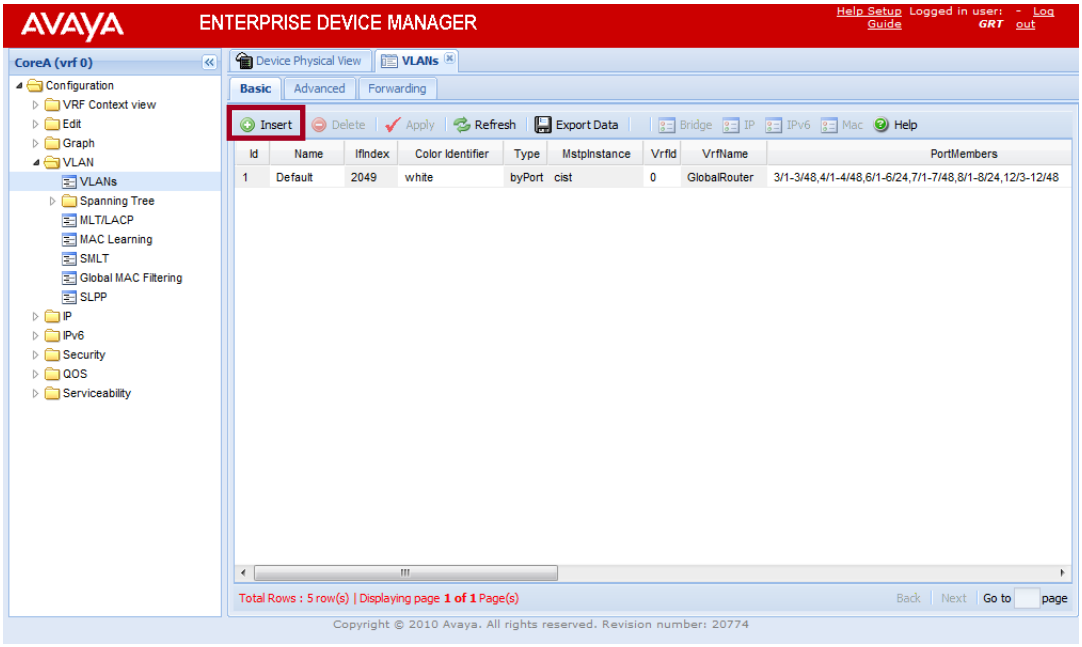

4 Set the *Id* to *100* then optionally specify a *Name*. Click '…' to assign *PortMembers*:

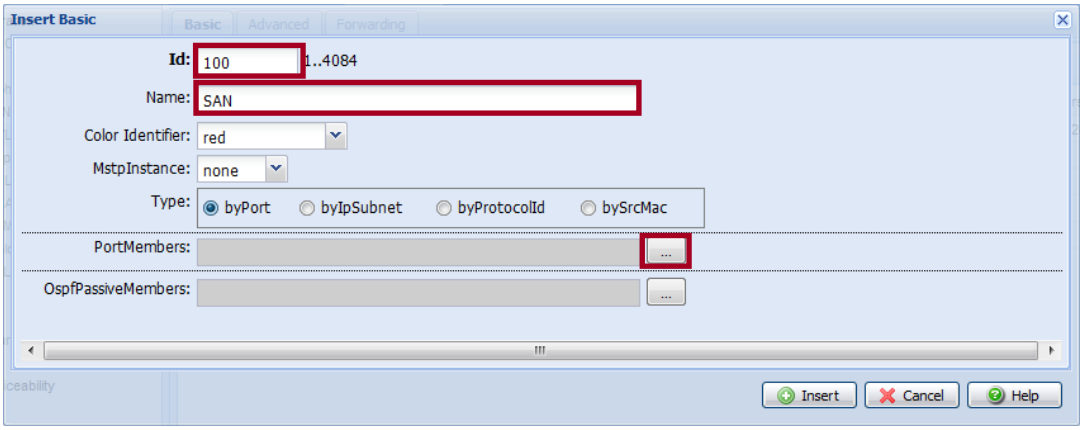

5 Select ports *3/1-3/9* and *6/1-6/9* then click *OK*:

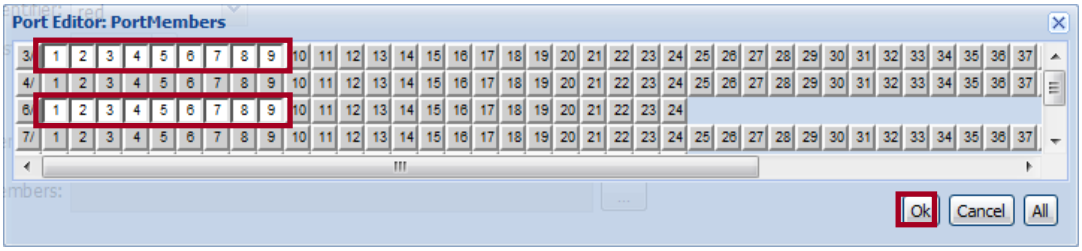

6 Click *Insert* to create the VLAN and apply the changes:

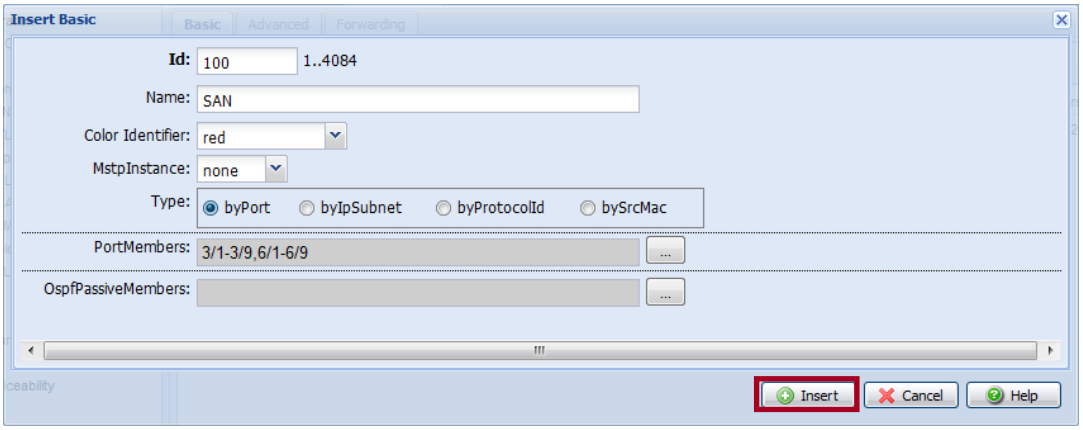

4.1.3Enabling Jumbo Frames

Enable jumbo frame support with the MTU set to at least 9600 bytes. To configure the MTU to 9600 bytes using the ACLI:

1 Create the port based VLAN 100:

```
VSP9000:1(config)# sys mtu 9600
```

Enable jumbo frame support with the MTU set to at least 9600 bytes. To configure the MTU to 9600 bytes using EDM:

1 Select *Configuration >Edit>Chassis > Chassis*. Set the *MTUSize* to *9600* then click *Apply*:

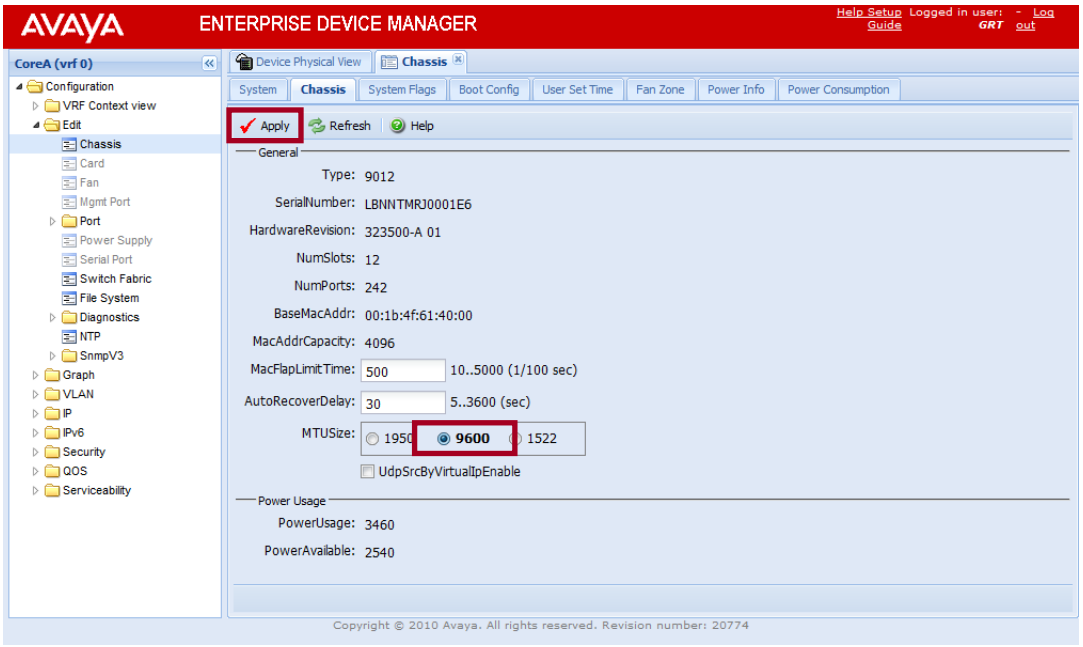

4.1.4Configuring Multiple Spanning Tree Protocol on the Avaya VSP 9000

By default MSTP is enabled globally on the Avaya VSP 9000 and is enabled on all ports. To enable fast convergence when a device is connected to the Avaya VSP 9000, the 10 GbE ports will be configured for edge mode.

To enable Admin Edge mode on ports 3/*1-3/9* and *6/1-6/9* using the ACLI:

1 Enter the configuration context for the Gigabit ports*3/1-3/9* and *6/1-6/9*:

```
VSP9000:1(config)# interface gigabitEthernet 3/1-3/9,6/1-6/9
```

2 Configure the ports for *MSTP Edge Mode*:

```
VSP9000:1(config-if)# spanning-tree mstp edge-port true
```

To enable Admin Edge mode on ports 3/*1-3/9* and *6/1-6/9* using the edm:

1 Select *Configuration >VLAN>Spanning Tree>MSTP > MSTI Port*. Toggle the *AdminEdgeStatus* for the ports to *true* then click *Apply*:

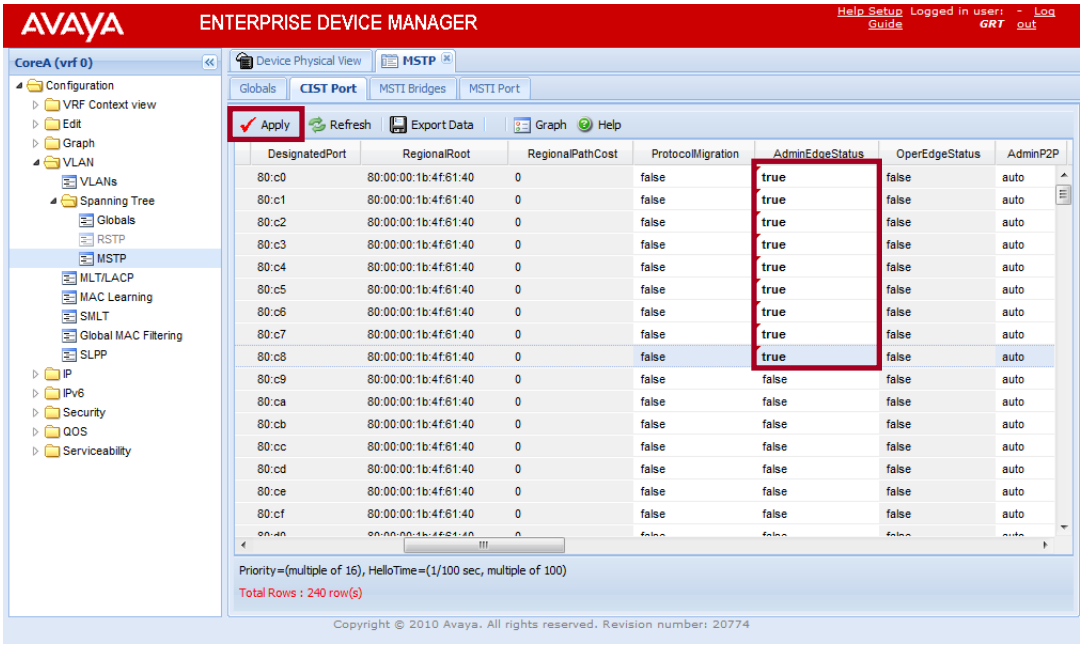

4.2 Avaya Virtual Services Platform 7000 Configuration

Follow the steps in this section to configure the Avaya VSP 7000. To allow for full utilization of the network for SAN storage, Coraid recommends isolating SAN storage networks from other network traffic.

Note – In this test setup, STP is disabled on all connected ports for faster transition time when a port is unplugged and plugged during. Also, the specific VLANs and ports used have no special significance and therefore can be changed at the user's discretion.

4.2.1 Configuring VLANs

Configure VLAN 100 and ensure that all ports connected to the Coraid HBAs and SRX-Series appliances are in VLAN 100. The following steps show how to configure the VLAN:

1 Enter the *Privilege Exec* context:

```
7024XLS>enable
```

2 Enter the *Global Configuration* context:

```
7024XLS# configure terminal
```

3 Create the port based VLAN 50:

```
7024XLS(config)# vlan create 100 name CORAID type port 1
```

4 Remove the ports from the default VLAN 1:

```
7024XLS(config)# vlan members remove 1 1-12
```

5 Add ports to the VLAN 50:

```
7024XLS(config)# vlan members add 100 1-12
```

4.2.2 Configuring Jumbo Frames, STP and the IP Address

Use the following steps to enable jumbo frame support, disable STP, and set the IP address:

1 In *Global Configuration* mode, enable jumbo frames:

```
7024XLS(config)# jumbo-frame enable
```

2 Confirm that jumbo frames are enabled:

```
7024XLS(config)# show jumbo-frames
```

3 Disable STP on the ports:

```
7024XLS(config)# interface fastEthernet 1-12
7024XLS(config-if)# spanning-tree learning disable
```

4 Set the IP address (for management purposes only):

```
7024XLS(config)# ip address 100.100.50.71 netmask 255.255.255.0 default-gateway
100.100.50.254
```

4.3 Server Configuration

This section describes how to configure the PC servers, which must run a 64-bit OS. The Avaya Networking Test Lab setup uses five PC servers running CentOS 5.5 (x64-bit version).

The Coraid HBA and associated driver must be installed in each PC server. The Coraid HBA requires a PCI Express slot and provides functionality similar to a standard NIC.

Before using the Coraid HBA SAN storage solution, one or more Logical Unit Adapters (LUN) must be configured on the Coraid EtherDrive SRX-Series storage appliance. More information on configuring Coraid appliances may be found at http://support.coraid.com/support/quickstart/. Figure 5 shows a sample LUN configuration:

1) create RAID groups
2) assign hot spares
3) bring LUNs online

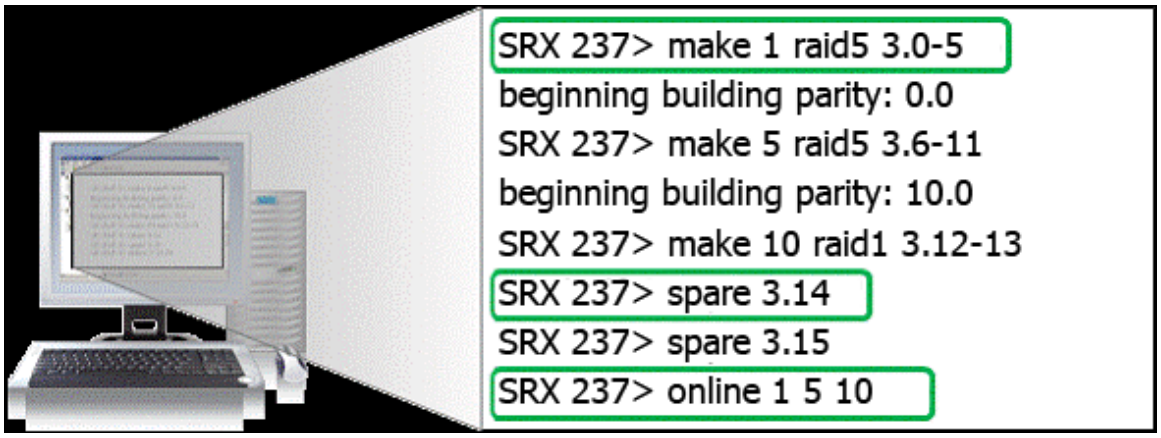

Figure 5 – Sample LUN Configuration

Use the following steps to configure the PC servers:

1) Use the ***make*** command to configure five different RAID groups of various sizes on the SRX-Series appliances. Coraid EtherDrive SRX-Series appliances automatically create one LUN per RAID group.
2) Use the ***spare*** command to assign hot spares.

3) Use the ***online*** command to bring the LUNs online.

4) Use the ***ethdrv-stat*** command to confirm that each PC can communicate with the drives and the drives are mountable.

5) Use the ***mkfs*** command to create and format a file system on each RAID. The AoE protocol makes the RAID look like a local storage device:
mkfs -t ext3 /dev/<device>

6) Each PC should mount a different RAID array:

```
mount -t ext3 /dev/<device> /mnt/coraid
```

7) Confirm that the RAID is reachable to the drive and write a file to it with the following command:

```
cd /mnt/coraid;touch foo
```

The Coraid HBAs have two ports and utilize every available port between the server initiator and target (see Figure 6). The HBAs have Coraid firmware to manage the Ethernet SAN and cannot be configured as a regular network card.

Coraid HBAs load balance and provide high performance link redundancy without the need for port aggregation, bonding, or traditional SAN MPIO software. Coraid EtherDrive SRX-Series storage appliances do not support MLT or LACP.

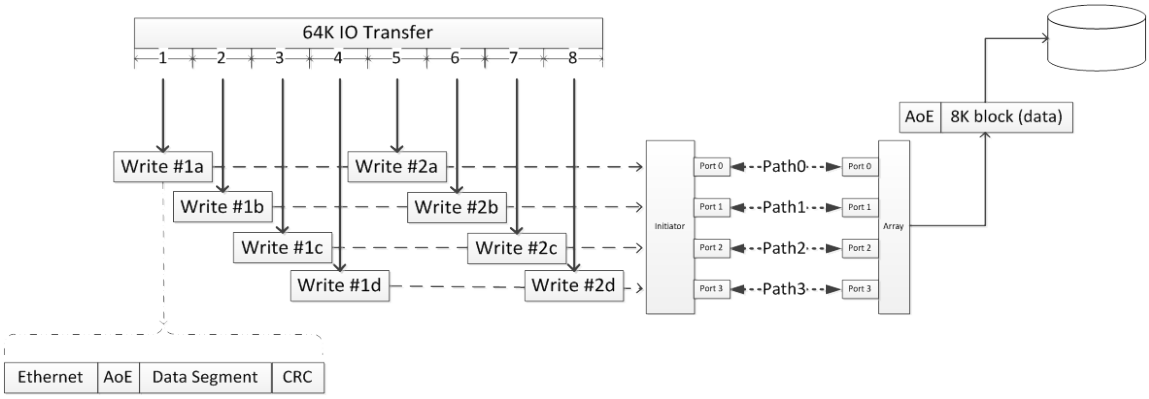

Figure 6 – Coraid Parallel Path Utilization w/o MPIO

For Coraid HBA support information, go to Coraid EtherDrive: Ethernet SAN for Virtualization, Cloud, and Enterprise Storage at http://www.coraid.com/support/customer_support.

4.4 EtherDrive SRX-Series Storage Appliance Configuration

In the Avaya Networking Test Lab setup, there are two Coraid EtherDrive SRX-Series storage appliances. Between the EtherDrive SRX appliances are five different RAID arrays so that each PC server writes to its own RAID. The RAID types are: RAID 1, RAID 5 and RAID 10 arrays.

For EtherDrive SRX-Series firmware and documentation support, go to http://www.coraid.com/support/customer_support.

Server	RAID Type	Number of Disks	Disk Type
1	1, 5 or 10	5	WD1002FBYS, 1TB, SATA, 7200 RPM
2	1, 5 or 10	5	WD1002FBYS, 1TB, SATA, 7200 RPM
3	1, 5 or 10	5	WD1002FBYS, 1TB, SATA, 7200 RPM
4	1, 5 or 10	5	WD1002FBYS, 1TB, SATA, 7200 RPM
5	1, 5 or 10	5	WD1002FBYS, 1TB, SATA, 7200 RPM

Table 2 – Storage Appliance Configuration

5. Test Methodology

The Avaya Networking Test Lab conducted comprehensive tests using a methodology that verified various fault conditions. Avaya's Switch Clustering uses the Split Multi-Link Trunking (SMLT) feature for Layer 2 Always-On Networking. A majority of the testing revolved around the use of the SMLT technology to ensure fast failover and uninterrupted access to the SAN devices.

Using Vdbench to write and then read x number of bytes, the Test Lab used the following scenarios to test the solution with a strong focus on High Availability (HA):

1) Test Case 1 – SMLT Baseline Test
2) Test Case 2 – Simulate SMLT Port Failure (ports 1 and 3 active)
3) Test Case 3 – Simulate SMLT Port Failure (ports 2 and 4 active)
4) Test Case 4 – Simulate SMLT Leg Failure (ports 1 and 2 active)
5) Test Case 5 – Simulate SMLT Leg Failure (ports 3 and 4 active)
6) Test Case 6 – Simulate SMLT Ports 1 and 3 Link Failover
7) Test Case 7 – Simulate SMLT Ports 2 and 4 Link Failover
8) Test Case 8 – Simulate SMLT Leg 1 Link Failover
9) Test Case 9 – Simulate SMLT Leg 2 Link Failover
10) Test Case 10 – Simulate broken SMLT/MLT (only one port active)
11) Test Case 11 – Simulate port Failure on SRX
12) Test Case 12 – Simulate port Failure on PC
13) Test Case 13 – Simulate SMLT-SRX Port Link Failover
14) Test Case 14 – Simulate SMLT-PC Port Link Failover

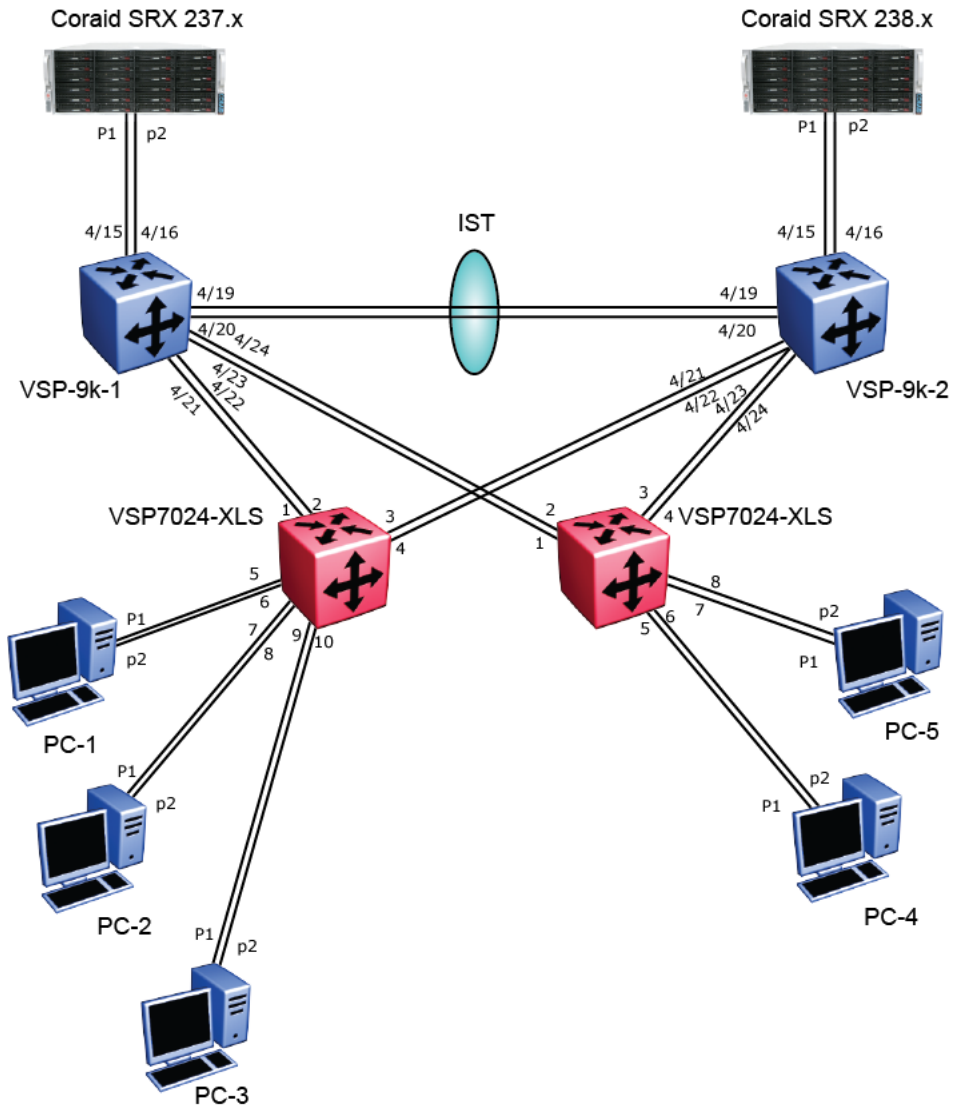

Figure 7 - Avaya Networking Test Lab Configuration

5.1 Test Case 1 – SMLT Baseline Test

The objective of this test case is to test the SMLT solution under normal conditions and retrieve baseline operating data.

Test Steps	Expected Results	Actual Results
1) Confirm that VLAN 100 was created on the VSP. All ports connected to the PC's HBAs and the SRXs should be in VLAN 100. Confirm that all ports come up.	All 10 GbE ports should have links and be up.	Pass
2) Configure five different RAID LUNs of various sizes on the SRX. After the LUNs are online, confirm that each PC can see the drives and the drives are mountable via ***ethdrv-stat***. Format each LUN for testing ***mkfs -t ext3 /dev/<device>***. Each PC should mount a different RAID array ***mount -t ext3 /dev/<device> /mnt/coraid***	Confirm that the RAID LUNs were created on the SRXs and can be seen from the PCs. From the PCs, you should be able to format the LUNs and then mount them.	Pass
3) Confirm that the LUN is reachable to the drive and write a file to it with the following command: ***cd /mnt/coraid;touch foo***	RAID LUNs should be readable and writable from the PCs. RAID LUNs should appear local to the PC.	Pass
4) Launch Vdbench on each PC with the target being the RAID LUN. Set the test to run for one hour, ***./vdbench -f param -vr***. This will read the parameter file for its configuration and do immediate data validation.	Vdbench should have a parameter file setup so that Vdbench knows which drive to write to and how much writing and reading it should do as well as duration. Vdbench should launch and begin reading and writing to the arrays.	Pass
5) Wait one hour for the reading and writing to complete. Vdbench stores all the results in html format in the output directory. Launch Firefox to read the results. Confirm that the error log contains no errors. Also confirm that the data verification confirms that all data was verified and accounted for.	The Vdbench error log should contain no errors. When Vdbench terminates, it should also verify that all data was written, read, and accurate.	Pass

5.2 Test Case 2 – Simulate SMLT Port Failure (ports 1 and 3 active)

The objective of this test case is to verify operations when one port from each leg of the SMLT fails.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should be flowing according to the MLT algorithm.	Pass
2) While Vdbench is running, send Layer 2 traffic to populate MAC tables on all switches. Remove one port from each leg of the SMLT (ports 2 and 4).	Check the VSP 7000/9000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX.		Pass

5.3 Test Case 3 – Simulate SMLT Port Failure (ports 2 and 4 active)

The objective of this test case is to verify operations when one port from each leg of the SMLT fails.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should be flowing according to the MLT algorithm.	Pass
2) While Vdbench is running, send Layer 2 traffic to populate MAC tables on all switches. Remove one port from each leg of the SMLT (ports 1 and 3).	Check the VSP 7000/9000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX.		Pass

5.4 Test Case 4 – Simulate SMLT Leg Failure (ports 1 and 2 active)

The objective of this test case is to verify operations when one leg of the SMLT fails.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should fail over to the active leg of the MLT.	Pass
2) While Vdbench is running, send Layer 2 traffic to populate MAC tables on all switches. Remove Leg 1 of the SMLT. (Ports 1 and 2 are the active ports.)	Check the VSP 7000/9000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX.		Pass

5.5 Test Case 5 – Simulate SMLT Leg Failure (ports 3 and 4 active)

The objective of this test case is to verify operations when one leg of the SMLT fails.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should fail over to the active leg of the MLT.	Pass
2) While Vdbench is running, send Layer 2 traffic to populate MAC tables on all switches. Remove Leg 2 of the SMLT. (Ports 3 and 4 are the active ports.)	Check the VSP 7000/9000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX.		Pass

5.6 Test Case 6 – Simulate SMLT Ports 1 and 3 Link Failover

The objective of this test case is to verify operations when SMLT ports 1 and 3 failover.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should fail over to the active port.	Pass
2) While Vdbench is running, send Layer 2 traffic to populate MAC tables on all switches. Remove port 1 and port 3. Then plug the ports back in to simulate link failover.	Check the VSP 7000/9000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX.		Pass

5.7 Test Case 7 – Simulate SMLT Ports 2 and 4 Link Failover

The objective of this test case is to verify operations when SMLT ports 2 and 4 failover.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should fail over to the active port.	Pass
2) While Vdbench is running, send Layer 2 traffic to populate MAC tables on all switches. Remove port 2 and port 4. Then plug the ports back in to simulate link failover.	Check the VSP 7000/9000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX.		Pass

5.8 Test Case 8 – Simulate SMLT Leg 1 Link Failover

The objective of this test case is to verify operations with SMLT leg 1 link failover.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should fail over to the active leg.	Pass
2) While Vdbench is running, send Layer 2 traffic to populate MAC tables on all switches. Remove leg 1 of the SMLT. Then plug leg 1 back in to simulate link failover.	Check the VSP 7000/9000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX.		Pass

5.9 Test Case 9 – Simulate SMLT Leg 2 Link Failover

The objective of this test case is to verify operations with SMLT leg 2 link failover.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should fail over to the active leg.	Pass
2) While Vdbench is running, send Layer 2 traffic to populate MAC tables on all switches. Remove leg 2 of the SMLT. Then plug leg 2 back in to simulate link failover.	Check the VSP 7000/9000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX.		Pass

5.10 Test Case 10 – Simulate Broken SMLT/MLT (only one port active)

The objective of this test case is to simulate a broken link where all ports failed except one.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should flow normally on the remaining port of the MLT.	Pass
2) While Vdbench is running, send Layer 2 traffic to populate MAC tables on all switches. Remove all ports except one (port 1 on the VSP 7000). This leaves only one port active on the SMLT/MLT link.	Check the VSP 7000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX.		Pass

5.11 Test Case 11 – Simulate Port Failure on SRX

The objective of this test case is to simulate port failure (on the SRX) while sending traffic and then validate that data continues to run on the remaining port.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should flow normally but at a slower throughput because one port is down.	Pass
2) Make sure SMLT/MLT is functioning. Remove port p1 from the SRX. Send Layer 2 traffic to populate MAC tables on all switches. Remove port p1 from each PC.	Check the VSP 7000/9000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX and MAC addresses are learned properly.		Pass

5.12 Test Case 12 – Simulate Port Failure on PC

The objective of this test case is to simulate port failure (on the PC) while sending traffic and then validate that data continues to run on the remaining port.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should flow normally but at a slower throughput because one port is down.	Pass
2) Make sure SMLT/MLT is functioning. Send Layer 2 traffic to populate MAC tables on all switches. Remove port p1 from each PC.	Check the VSP 7000/9000 MAC table and port statistics to make sure traffic is forwarding and MAC is learning properly.	Pass
3) Verify that I/O continues from PC to SRX and MAC addresses are learned properly.		Pass

5.13 Test Case 13 – Simulate SMLT-SRX Port Link Failover

The objective of this test case is to simulate an SRX port link failover.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should flow normally but at a slower throughput because one port is down.	Pass
2) Make sure SMLT/MLT is functioning. Remove port p1 from the SRX. Wait 2 minutes and plug the port back in.	Coraid driver takes 80-100 seconds to detect and mark the link as down. After the link is plugged back in, it takes 8-10 secs for I/O traffic to resume on that link. The downed link that came back up had to re-initialize and do SAN disk discovery before I/O traffic starts flowing again. Test observation – If you unplug a link and plug it back in quickly (1-2 secs), I/O traffic resumes on that link sub-second since the link is not yet marked down by the Coraid driver.	Pass
3) Verify that I/O continues from PC to SRX.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

5.14 Test Case 14 – Simulate SMLT-PC Port Link Failover

The objective of this test case is to simulate a PC port link failover.

Test Steps	Expected Results	Actual Results
1) Start Vdbench.	I/O traffic should flow normally but at a slower throughput because of the port flip flops.	Pass
2) Make sure SMLT/MLT is functioning. Remove one port p2 from each PC. Wait 2 minutes and plug the ports back in.	Coraid driver takes 80-100 seconds to detect and mark the link as down. After the link is plugged back in, it takes 8-10 secs for I/O traffic to resume on that link. The downed link that came back up had to re-initialize and do SAN disk discovery before I/O traffic starts flowing again. Test observation – If you unplug a link and plug it back in quickly (1-2 secs), I/O traffic resumes on that link sub-second since the link is not yet marked down by the Coraid driver.	Pass
3) Verify that I/O continues from PC to SRX.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

6. Conclusion

The Avaya Networking Test Lab conducted comprehensive tests using a methodology that verified throughput under normal conditions and under various fault conditions. There were no errors observed during the testing of the various fault conditions.

The test lab confirmed through various scenarios that the solution passed all High Availability (HA) tests. The test lab simulated many HA fault conditions to ensure that the PCs continued to read and write to the storage arrays with no errors. After thorough interoperability testing, all test cases were successfully completed and confirm that this solution meets Avaya's quality and interoperability standards.

7. Additional Resources

- For Avaya product documentation, go to http://support.avaya.com/.
- For Coraid product documentation, go to http://www.coraid.com/products.
- For Coraid HBA support information, go to http://www.coraid.com/support/customer_support.
- For Coraid configuration information, go to http://support.coraid.com/support/quickstart/.
- For Coraid SRX firmware and documentation support, go to http://www.coraid.com/support/customer_support.
- For AoE information, go to http://www.coraid.com/support/faqs.
- For Vdbench downloads and documentation, go to http://sourceforge.net/projects/vdbench/.