Virtual Services Platform 9000

Engineering

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

Avaya Data Solutions

Document Date: December 2010

Document Number: NN48500-611

Document Version: 1.1

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

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

Revision Control

No	Date	Version	Revised By	Remarks
1	December 2010	1.0	B. Azimi	Initial Draft

Table of Contents

Figures 6
Tables 7
1. Introduction 9
1.1 Interoperability testing 10
2. Reference configuration 11
3. Equipment and software validated 12
4. Storage area network configuration 13
4.1 Virtual Services Platform 9000 configuration 13
4.1.1 Configuring VLANs using the ACLI 13
4.1.2 Configuring VLANs using EDM 14
4.1.3 Enable jumbo frames 16
4.1.4 Multiple Spanning Tree Protocol 17
4.2 Server configuration 18
4.3 EtherDrive SRX-Series storage appliance configuration 20
5. Test methodology 21
5.1 Test Case 1 – Baseline test 22
5.2 Test Case 2 – Individual HBA link failure 23
5.3 Test Case 3 – Restore the failed HBA link 23
5.4 Test Case 4 – Simulate module failure by removing module 3 24
5.5 Test Case 5 – Reinsert the failed module 3 24
5.6 Test Case 6 – Simulate module failure by disabling the module 25
5.7 Test Case 7 – Restore the failed module with the ACLI 25
5.8 Test Case 8 – Simulate module failure by removing module 6 26
5.9 Test Case 9 – Reinsert the failed module 6 26
5.10 Test Case 10 – Simulate primary CP failure to test HA 27
5.11 Test Case 11 – Reinsert the failed primary CP 27
5.12 Test Case 12 – Simulate primary CP failure to test failover 28
5.13 Test Case 13 –Simulate HBA failure 29
6. Test results 30
7. Conclusion 33
8. Additional Resources 34

Figures

Figure 1 – SAN test bed configuration 11
Figure 2 – Sample LUN Configuration 18
Figure 3 – Coraid parallel path utilization w/o MPIO 19
Figure 3 – Vdbench results for test cases 1 to 6 30
Figure 4 – Vdbench results for test cases 7 to 12 31
Figure 5 – EDM port statistics (screen 1) 31
Figure 6 – EDM port statistics (screen 2) 32

Tables

Table 1 – Hardware and Software Versions 12
Table 2 – Storage Appliance Configuration 20

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 (3.0.0.0_B325) (PRIVATE)
        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 a solution comprised of the Avaya Virtual Services Platform 9000 (VSP 9000) and the Coraid EtherDrive™ SRX4200 Ethernet SAN array.

The VSP 9000 is a new Ethernet Switching platform for Super Large Campus environments and Enterprise Data Centers. This platform offers an unmatched architecture that scales from an initial 8.4 Terabits per second to an Industry-leading 27 Terabits per second. The VSP 9000 delivers industry-leading 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 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 without impacting user applications
- Layer 2 and 3 network virtualization services providing support for multiple customers and user groups on the same platform
- Network failover in less than 20 milliseconds
- 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
- Shortest Path Bridging – de-couples infrastructure from connectivity services and ensures predictability for all network services

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.

1.1 Interoperability testing

The interoperability testing focused primarily on verifying throughput under normal conditions and then simulating various fault conditions to test high availability. During interoperability testing, the VSP 9000 communicated with two Coraid EtherDrive SRX-Series storage appliances.

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 are block storage RAID appliances with front loading, hot-swappable SATA, SAS, and SSD disk drives. You can use each EtherDrive SRX appliance individually or in RAID sets.

The Coraid EtherDrive SRX-Series use 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. Ethernet 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 no data loss is guaranteed.

For technical support on Coraid EtherDrive SRX-Series storage appliances, contact Coraid's technical support at their Web site:

http://www.coraid.com/support/customer_support

The Web site provides documentation, frequently asked questions, AoE tools for Linux, and contact telephone numbers.

2. Reference configuration

Figure 1 shows a sample storage area network (SAN) configuration with an VSP 9000, five PC servers and two Coraid EtherDrive SRX-Series storage appliances. All communication between these devices is over 10 GbE links in a layer two 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 and measure the performance of the Coraid EtherDrive SRX-Series storage appliances, the Avaya Data Solutions Test Lab used the Open Source software, Vdbench. Vdbench 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/.

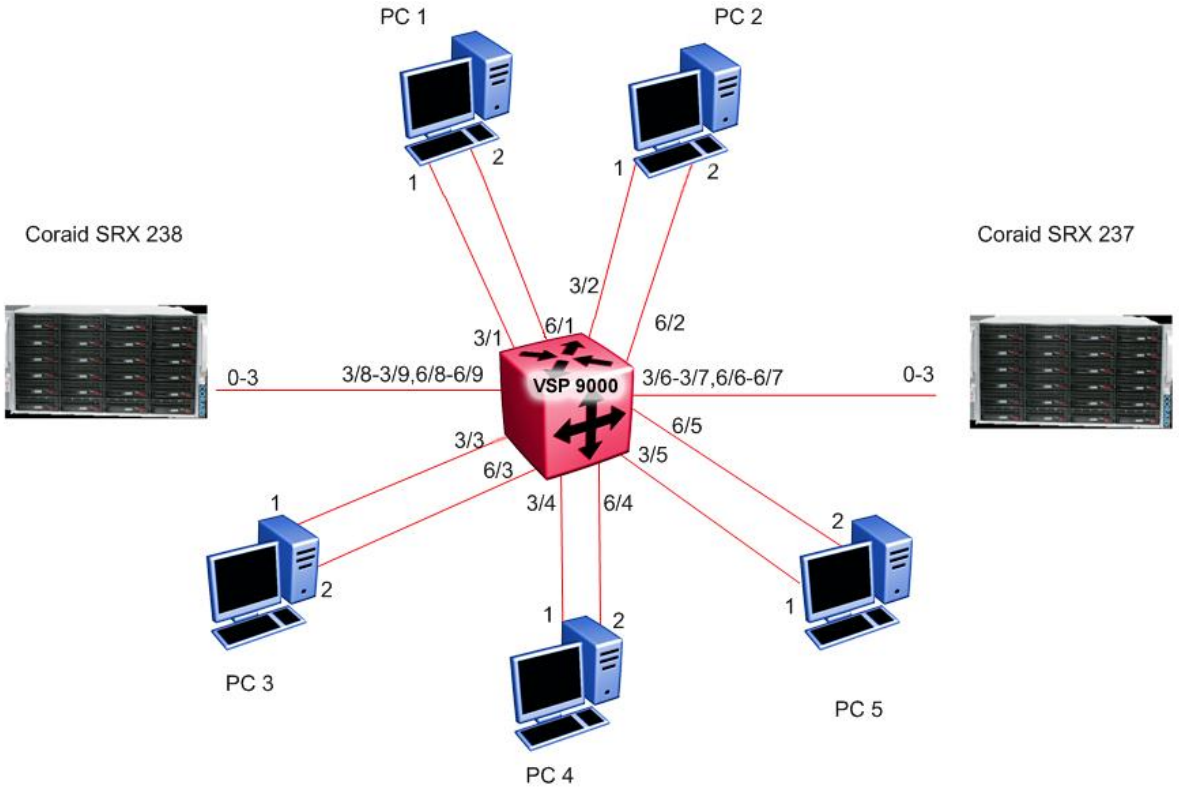

Figure 1 – SAN test bed configuration

3. Equipment and software validated

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

Equipment	Software
• 1 x VSP 9000	Software Release 3.0.0.0b327
• 5 x PC 1U servers with Coraid Host Bus Adapter (HBA) network cards installed	Coraid HBA build 1.3.1.1-00051 CentOS 5.5 64-bit
• 2 x SRX-Series storage appliances Each SRX-Series appliance has: • 2 x two-port 10Gb/s network cards installed. • 25 x 300GB SATA drives	CorOS 5.1.4
• 5 x Coraid 10G, dual port, SFP+ HBA Cards (1 per server)	
	Vdbench verification software

Table 1 – Hardware and Software Versions

4. Storage area network configuration

This section describes how to set up the following SAN components:

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

4.1 Virtual Services Platform 9000 configuration

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

Note – The specific VLANS and ports used in this test setup have no special significance. You can use whatever VLANs and ports you want.

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
```

4.1.2 Configuring VLANs using EDM

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 EDM:

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

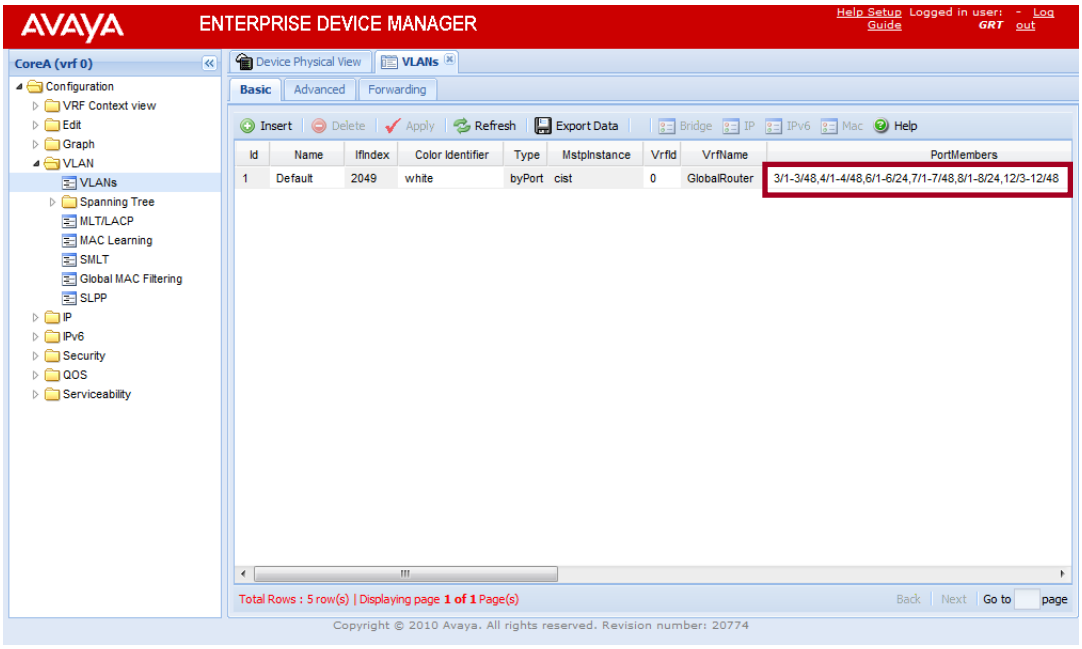

2 Uncheck ports *1/1-1/9* and *6/1-6/9* then click *OK*:

3 Create a new VLAN by clicking *Insert*:

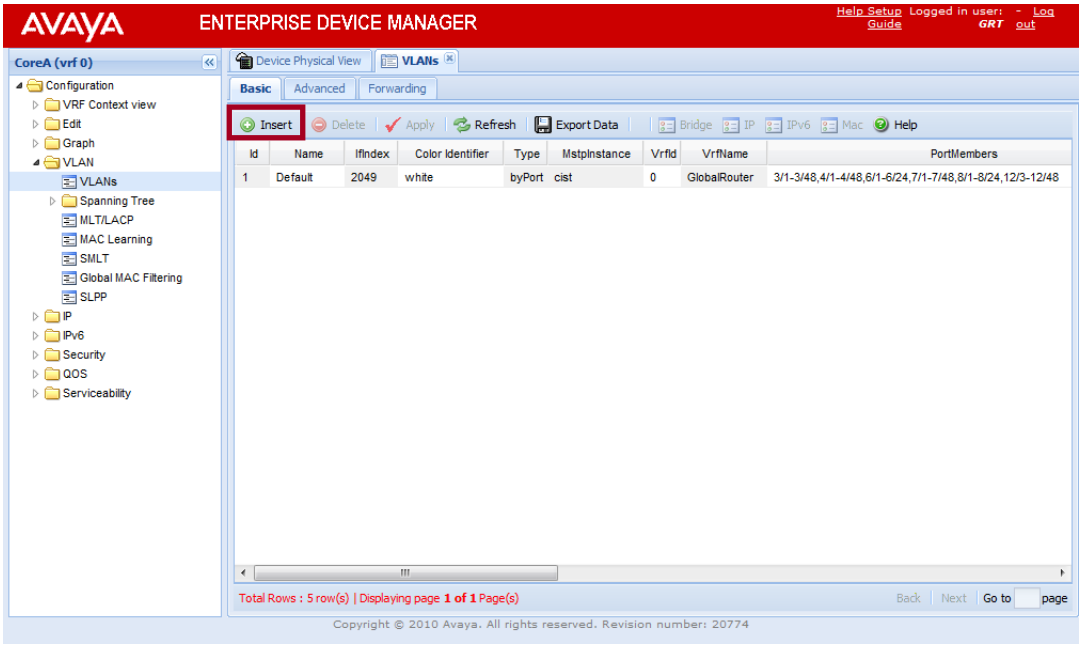

4 Set the *Id* to *100* then optionally specify a *Name*. Click '…' to assign *PortMembers*:

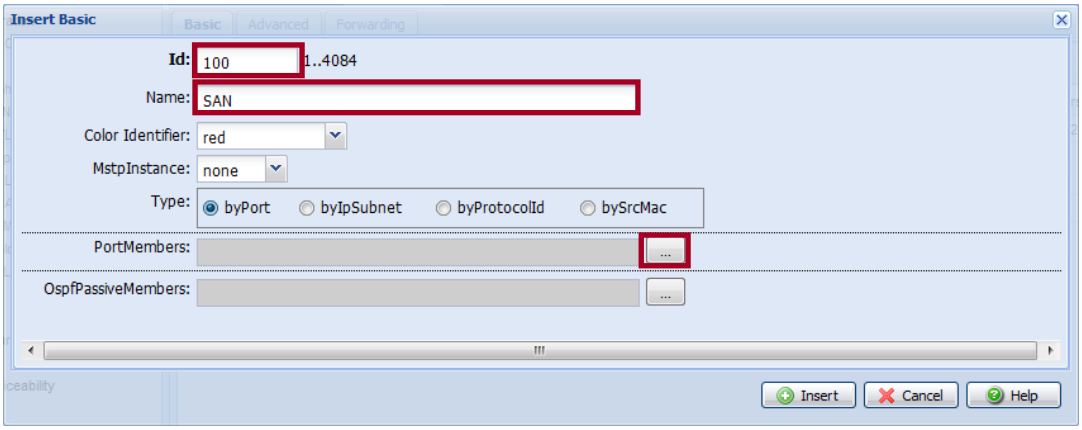

5 Select ports *3/1-3/9* and *6/1-6/9* then click *OK*:

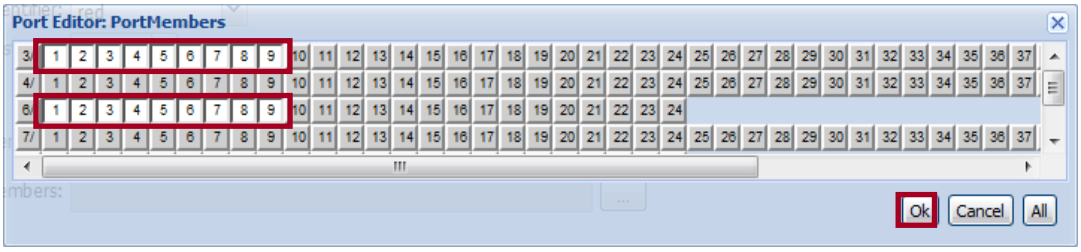

6 Click *Insert* to create the VLAN and apply the changes:

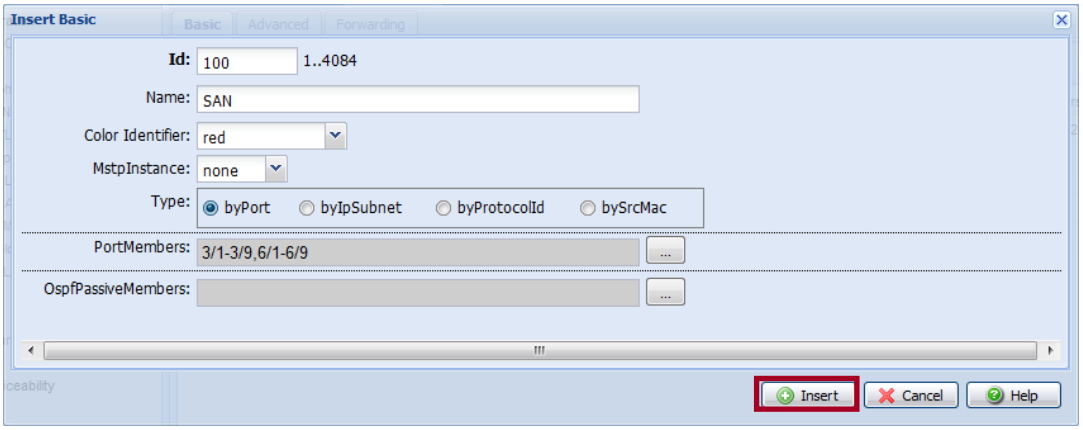

4.1.3 Enable jumbo frames

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 supportwith 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*:

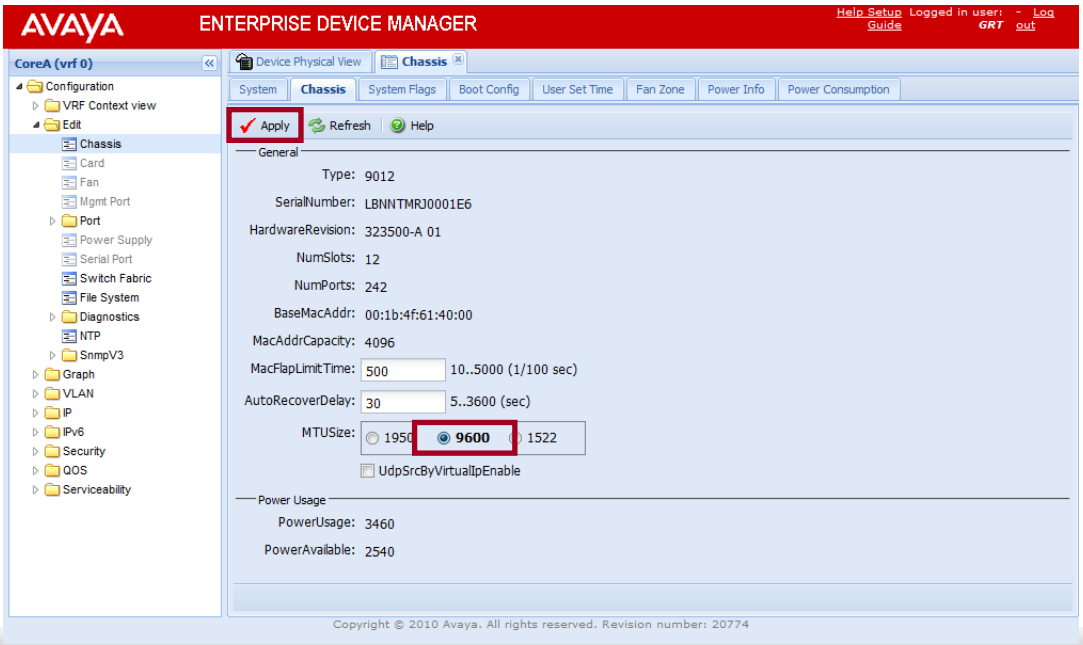

4.1.4 Multiple Spanning Tree Protocol

By default MSTP is enabled globally on the VSP 9000 and is enabled on all ports. To enable fast convergence when a device is connected to the 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*:

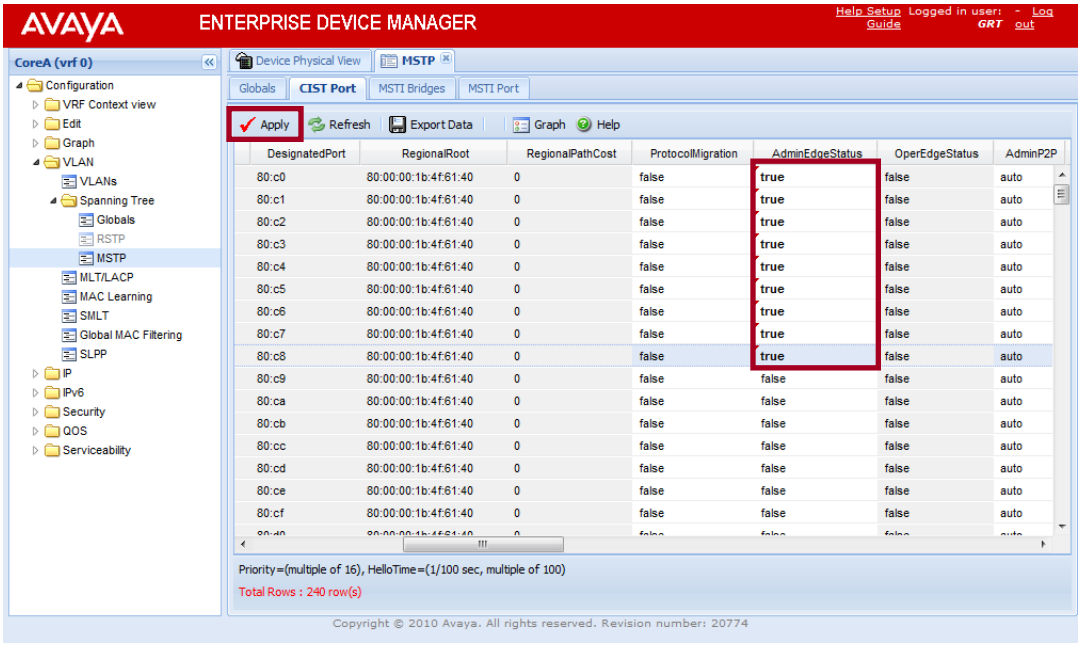

4.2 Server configuration

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

You also have to install Coraid HBA cards in each PC server. The Coraid HBA provides functionality similar to a NIC card. The Coraid HBA card requires a PCI Express slot. After installing the Coraid HBA cards, you also have to install the Coraid HBA driver in each PC server.

Before you can use the SAN storage through your Coraid HBA, you have to configure one or more Logical Unit Adapters (LUN) on the Coraid EtherDrive SRX-Series storage appliance. Figure 2 shows a sample LUN configuration where with just three commands (make, spare and online), you can:

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

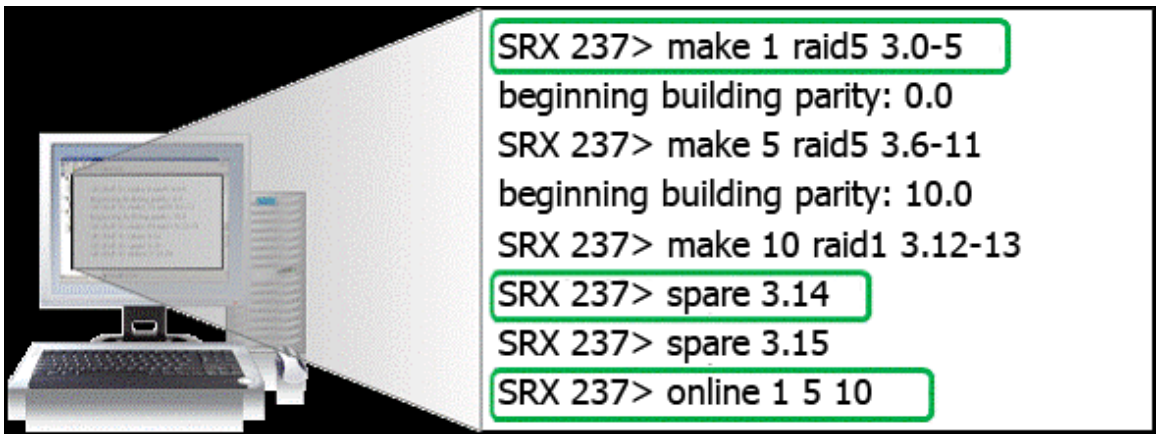

Figure 2 – Sample LUN Configuration

Use the following steps to configure the PC servers:

1) Use the ***make*** command to configure five different RAIDS 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 RAIDs 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 3). The HBAs have Coraid firmware to manage the Ethernet SAN and cannot be configured as a regular network card.

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

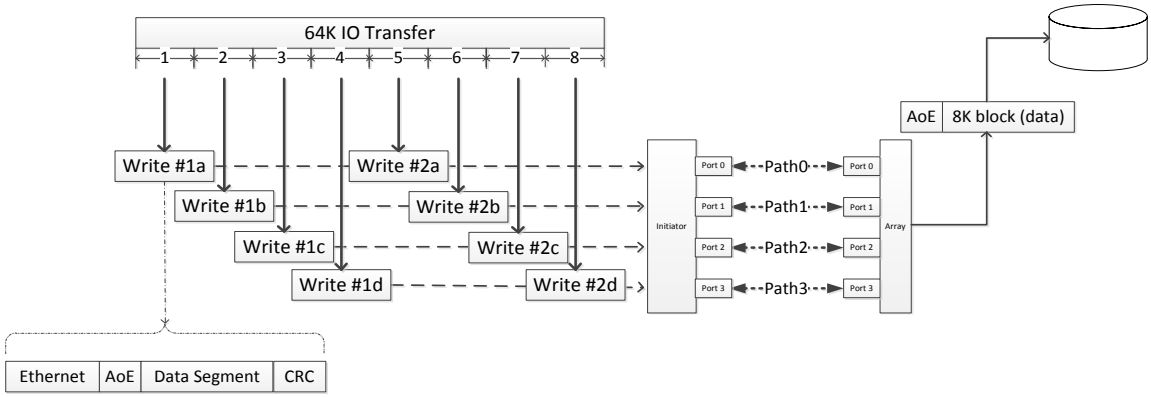

Figure 3 – Coraid parallel path utilization w/o MPIO

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.

4.3 EtherDrive SRX-Series storage appliance configuration

In the Avaya Data Solutions 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 Data Solutions Test Lab conducted comprehensive tests using a methodology that verified throughput under normal conditions and under various fault conditions. 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 – Baseline test of the solution.
2) Test Case 2 – Simulate HBA link failure.
3) Test Case 3 – Restore HBA link.
4) Test Case 4 – Simulate VSP 9000 module 3 failure by removing module.
5) Test Case 5 – Restore VSP 9000 module by replacing module 3.
6) Test Case 6 – Simulate VSP 9000 module 3 failure by disabling module with ACLI.
7) Test Case 7 – Restore VSP 9000 module by re-enabling module 3 with ACLI.
8) Test Case 8 – Simulate VSP 9000 module 6 failure by removing module.
9) Test Case 9 – Restore VSP 9000 module by replacing module 6.
10) Test Case 10 – Enable HA on the VSP and then remove the primary CP to test data loss.
11) Test Case 11 – Re-install the CP removed in Scenario 10.
12) Test Case 12 – Enable HA on the VSP and then remove the primary CP to test failover.
13) Test Case 13 – Simulate HBA failure by removing both links and then re-inserting to test recovery.

5.1 Test Case 1 – Baseline test

The objective of this test case is to test the 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 link and be up.	Pass
2) Configure five different RAIDS of various sizes on the SRX. After the RAIDs are online, confirm that each PC can see the drives and the drives are mountable ***ethdrv-stat***. Format each RAID 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 arrays were created on the SRXs and can be seen from the PCs. From the PCs, you should be able to format the RAIDs and then mount them.	Pass
3) Confirm that the RAID is reachable to the drive and write a file to it with the following command: ***cd /mnt/coraid;touch foo***	RAID arrays should be readable and writable from the PCs. RAID arrays should appear local to the PC.	Pass
4) Launch Vdbench on each PC with the target being the RAID array. 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 – Individual HBA link failure

The objective of this test case is to verify what happens when one link fails.

Test Steps	Expected Results	Actual Results
1) The HBAs have two 10 GbE links. Confirm that one link is on module three of the VSP and the other link is on module six of the VSP.	Both links on the HBAs should be up. The command **`ethdrv-stat`** should show the RAIDs reachable from both HBA ports. Confirm on the VSP as well that the links are up.	Pass
2) The HBAs will use any available port to read and write data. Start Vdbench on all the PCs and confirm that the tool is working for each.	All PCs should be reading and writing to the storage array.	Pass
3) With traffic running, from one PC, unplug one link from either the HBA or from the VSP. Look at the VSP KHI stats for both HBA links to see which one has the greater amount of traffic and unplug that link.	Vdbench should keep on reading and writing to the array as if nothing happened.	Pass
4) Wait for the test to finish to see if any data was lost. Check both the output at the end of Vdbench as well as the summary and error logs on.	Vdbench should not have any errors or invalid data issues so that you know that all data was correctly processed.	Pass

5.3 Test Case 3 – Restore the failed HBA link

The objective of this test case is to restore the failed link from Test Case 2 to validate that data resumes running on both links.

Test Steps	Expected Results	Actual Results
1) Run Vdbench from all the PCs. Confirm that Vdbench is running correctly.	Vdbench should be reading and writing on all PCs.	Pass
2) On the PC that only has one link, wait a few minutes after Vdbench has been running and then plug in the secondary link.	Vdbench should continue to read and write with no issues.	Pass
3) Wait for Vdbench to finish testing. Confirm that all data is validated and the error log has no errors.	Vdbench should validate all data and the error log in the output directory should be empty.	Pass

5.4 Test Case 4 – Simulate module failure by removing module 3

The objective of this test case is to simulate a VSP 9000 module failure by removing module 3 and then validating that each PC continues to read and write to their RAID.

Test Steps	Expected Results	Actual Results
1) Confirm that all links are up on the VSP. Start Vdbench from the PCs.	Vdbench should launch and begin reading and writing to the storage array.	Pass
2) After Vdbench has been running for a few minutes, remove module 3 from the VSP.	Since half the links are on module 3, and the other half are on module 6, we should be able to remove an I/O module and all devices will still have a link and be reachable.	Pass
3) Confirm that Vdbench on each PC continues reading and writing to the storage array.	Confirm that Vdbench continues to read and write to the storage array with half the links available.	Pass
4) Wait for Vdbench to finish reading and writing to verify the results.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

5.5 Test Case 5 – Reinsert the failed module 3

The objective of this test case is to reinsert module 3 from Test Case 4 and then validate that data resumes running on the failed module.

Test Steps	Expected Results	Actual Results
1) Restart Vdbench again with only one link per HBA card and two links per SRX. Module 3 on the VSP should still be removed.	Vdbench should start with no issues.	Pass
2) After Vdbench has been running for a few minutes, reinsert module 3 in the VSP.	Module 3 should come up and the links going to the HBAs and SRXs should come up with no issues. Vdbench should continue to read and write to the storage array.	Pass
3) Wait for Vdbench to finish reading and writing to verify the results.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

5.6 Test Case 6 – Simulate module failure by disabling the module

The objective of this test case is to simulate VSP 9000 module failure by disabling module 3 with the ACLI and then validating that each PC continues to read and write to their RAID.

Test Steps	Expected Results	Actual Results
1) Confirm that all links are up on the VSP. Start Vdbench from the PCs.	Vdbench should launch and begin reading and writing to the storage array.	Pass
2) After Vdbench has been running for a few minutes, disable module 3 from the VSP.	Since half the links are on module 3, and the other half are on module 6, we should be able to disable an I/O module and all devices will still have a link and be reachable.	Pass
3) Confirm that Vdbench on each PC continues reading and writing to the storage array.	Confirm that Vdbench continues to read and write to the storage array with half the links available.	Pass
4) Wait for Vdbench to finish reading and writing to verify the results.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

5.7 Test Case 7 – Restore the failed module with the ACLI

The objective of this test case is to restore the failed module from Test Case 6 and then validate that data resumes running on the failed module.

Test Steps	Expected Results	Actual Results
1) Restart Vdbench again with only one link per HBA card and two links per SRX. Module 3 on the VSP should still be removed.	Vdbench should start with no issues.	Pass
2) After Vdbench has been running for a few minutes, enablemodule 3 using the command line on the VSP.	Module3 should come up and the links going to the HBAs and SRXs should come up with no issues. Vdbench should continue to read and write to the storage array.	Pass
3) Wait for Vdbench to finish reading and writing to verify the results.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

5.8 Test Case 8 – Simulate module failure by removing module 6

The objective of this test case is to simulate VSP 9000 module failure by removing module 6 and then validating that each PC continues to read and write to their RAID.

Test Steps	Expected Results	Actual Results
1) Confirm that all links are up on the VSP. Start Vdbench from the PCs.	Vdbench should launch and begin reading and writing to the storage array.	Pass
2) After Vdbench has been running for a few minutes, remove module6 from the VSP.	Since half the links are on module6, and the other half are on module3, we should be able to remove an I/O module and all devices will still have a link and be reachable.	Pass
3) Confirm that Vdbench on each PC continues reading and writing to the storage array.	Confirm that Vdbench continues to read and write to the storage array with half the links available.	Pass
4) Wait for Vdbench to finish reading and writing to verify the results.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

5.9 Test Case 9 – Reinsert the failed module 6

The objective of this test case is to reinsert the failed module from Test Case 8 and then validate that data resumes running on the failed module.

Test Steps	Expected Results	Actual Results
1) Restart Vdbench again with only one link per HBA card and two links per SRX. Module6 on the VSP should still be removed.	Vdbench should start with no issues.	Pass
2) After Vdbench has been running for a few minutes, reinsert module6 in the VSP.	Module 6 should come up and the links going to the HBAs and SRXs should come up with no issues. Vdbench should continue to read and write to the storage array.	Pass
3) Wait for Vdbench to finish reading and writing to verify the results.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

5.10 Test Case 10 – Simulate primary CP failure to test HA

The objective of this test case is to test HA by confirming that removing the master CP causes no data loss.

Test Steps	Expected Results	Actual Results
1) Run Vdbench from all the PCs. Confirm that Vdbench is running correctly.	Vdbench should be reading and writing on all PCs.	Pass
2) Ensure that HA is enabled on the VSP with **`show boot config flags`**.	HA mode should be enabled on the VSP.	Pass
3) Remove the master CP on the VSP 9000, the standby CP should come up as the master.	Standby CP should come up as master. Confirm that Vdbench doesn't notice the failure and continues reading and writing data to the storage array.	Pass
4) Wait for Vdbench to finish reading and writing to verify the results.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

5.11 Test Case 11 – Reinsert the failed primary CP

The objective of this test case is to reinsert the failed primary CP from Test Case 10 to validate that the CP comes up as the standby CP.

Test Steps	Expected Results	Actual Results
Run Vdbench from all the PCs. Confirm that Vdbench is running correctly	Vdbench should be reading and writing on all PCs.	Pass
Ensure that HA is enabled on the VSP with ***show boot config flags***.	HA mode should be enabled on the VSP.	Pass
Insert the CP that was removed in TC 10. This CP should come up as the standby.	The newly inserted CP should come up as a standby CP and the master CP should be unaffected. From the console, confirm that no error messages appear.	Pass
Wait for Vdbench to finish reading and writing to verify the results.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

5.12 Test Case 12 – Simulate primary CP failure to test failover

The objective of this test case is to test failover by confirming that removing the master CP causes no data loss.

Test Steps	Expected Results	Actual Results
1) Run Vdbench from all the PCs. Confirm that Vdbench is running correctly.	Vdbench should be reading and writing on all PCs.	Pass
2) Ensure that HA is enabled on the VSP with ***show boot config flags***.	HA mode should be enabled on the VSP.	Pass
3) Remove the master CP on the VSP 9000, the standby CP should come up as the master.	Standby CP should come up as master. Confirm that Vdbench doesn't notice the failure and continues reading and writing data to the storage array.	Pass
4) Wait for Vdbench to finish reading and writing to verify the results.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass

5.13 Test Case 13 –Simulate HBA failure

The objective of this test case is to test HBA recovery by removing both links from one HBA for 20 seconds and then plugging them back in to see if the PC recovers.

Test Steps	Expected Results	Actual Results
1) Run Vdbench from all the PCs again. Confirm that Vdbench is running correctly.	Vdbench should be reading and writing on all PCs.	Pass
2) Remove both links of an HBA for 20 seconds, either from the HBA itself or from the VSP.	Vdbench should hang Coraid states a failover of 15 seconds, so if the links are out for 20 seconds, we should lose traffic.	Pass
3) Plug both HBAs back in. Verify if the PC is stable and traffic is transmitting again.	Because the RAID arrays are mounted to the PCs, if the RAID is down for too long, the PC acts like it has lost a disk drive. This can cause the PC to hang. Verify that within 20 seconds a PC can lose the connection but recover.	Pass
4) Wait for Vdbench to finish reading and writing to verify the results.	Vdbench should finish with no errors in the log file and all data should be validated.	Pass
5) Multiple failover times for this scenario might need to be tested to see what the threshold is for a complete failover and the system can recover.	Try multiple dual link failure times to verify if a certain time can be established as the threshold where the system won't recover.	Pass

6. Test results

This Technical Configuration Guide describes a solution comprised of the VSP 9000 and Coraid storage appliances. After very thorough interoperability testing, all test cases were successfully completed. The following screens show the testing statistics, which confirm that this solution meets Avaya's quality and interoperability standards.

Figures 3 and 4 show the Vdbench statistics, which provides throughput data and validates that there were no errors:

Test Case	PC	#intervals	i/o rate Avg.	MB/sec Avg.	bytes i/o Avg.	read pct Avg.	resp time Avg.	resp max Avg.	resp stddev Avg.	Errors:
1	3	600	42.66	42.66	1048576	66.4	186.906	3670.882	301.573	None
	4	600	70.65	70.65	1048576	66.59	112.454	554.001	53.123	None
	5	600	38.33	38.33	1048576	66.43	208.063	3728.764	315.133	None
	6	600	2.74	27.43	10485760	66.1	2908.347	28349.462	3676.123	None
2	3	600	33.24	33.24	1048576	66.37	240.196	4122.832	412.368	None
	4	600	70.44	70.44	1048576	66.58	112.801	566.518	53.109	None
	5	600	31.95	31.95	1048576	66.38	249.795	3256.647	412.987	None
	6	600	4.08	40.82	10485760	66.67	1951.647	10350.287	926.393	None
3	3	600	32.19	32.19	1048576	66.36	247.781	8922.056	459.259	None
	4	600	70.65	70.65	1048576	66.59	112.466	536.646	52.941	None
	5	600	30.2	30.2	1048576	66.31	264.354	9381.611	480.266	None
	6	600	4.07	40.65	10485760	66.61	1956.667	6887.305	924.795	None
4	3	600	50.21	50.21	1048576	66.41	158.671	112136.555	1274.665	None
	4	600	74.73	74.73	1048576	66.61	106.272	112287.063	1086.221	None
	5	600	41.08	41.08	1048576	66.39	194.013	84219.8	1098.334	None
	6	600	3.76	37.65	10485760	66.3	2133.163	92425.757	5115.467	None
5	3	600	45.07	45.07	1048576	66.42	176.739	3154.615	162.539	None
	4	600	72.1	72.1	1048576	66.6	110.18	566.431	51.428	None
	5	600	40.95	40.95	1048576	66.42	194.541	2760.323	168.421	None
	6	600	3.98	39.78	10485760	66.55	2000.427	8028.367	939.384	None
6	3	600	48.84	48.84	1048576	66.43	163.027	84160.702	1189.588	None
	4	600	64.32	64.32	1048576	66.58	123.597	112321.292	1090.364	None
	5	600	37.56	37.56	1048576	66.39	212.204	84119.221	1697.18	None
	6	600	3.83	38.31	10485760	66.41	2080.022	5629.897	813.487	None

Figure 3 – Vdbench results for test cases 1 to 6

Test Case	PC	#intervals	i/o rate Avg.	MB/sec Avg.	bytes i/o Avg.	read pct Avg.	resp time Avg.	resp max Avg.	resp stddev Avg.	Errors:
7	3	600	33.41	33.41	1048576	66.36	238.674	3768.17	408.346	None
	4	600	71.01	71.01	1048576	66.57	111.899	502.124	53.024	None
	5	600	32	32	1048576	66.37	249.245	3566.278	414.369	None
	6	600	4.11	41.09	10485760	66.56	1930.061	6600.882	869.878	None
8	3	600	50.12	50.12	1048576	66.41	158.946	84195.157	1265.961	None
	4	600	72.83	72.83	1048576	66.59	109.058	57308.587	688.155	None
	5	600	41.35	41.35	1048576	66.4	192.805	84134.956	1128.684	None
	6	600	3.57	35.71	10485760	66.34	2232.456	84473.329	4769.588	None
9	3	600	37.26	37.26	1048576	66.34	213.919	10973.424	415.853	None
	4	600	72.31	72.31	1048576	66.59	109.862	543.48	51.459	None
	5	600	34.18	34.18	1048576	66.42	233.305	9422.087	424.245	None
	6	600	3.99	39.88	10485760	66.35	1995.463	11009.597	1022.319	None
10	3	600	32.74	32.74	1048576	66.36	243.72	4701.172	434.152	None
	4	600	70.83	70.83	1048576	66.6	112.18	671.405	53.4	None
	5	600	30.64	30.64	1048576	66.3	260.582	6430.106	455.458	None
	6	600	4.05	40.52	10485760	66.63	1969.446	7162.501	922.605	None
11	3	600	31.76	31.76	1048576	66.34	251.394	4960.547	442.384	None
	4	600	62.17	62.17	1048576	66.56	127.891	762.872	85.272	None
	5	600	30.55	30.55	1048576	66.31	260.959	5765.32	449.786	None
	6	600	41.24	41.24	10485760	66.44	1931.331	8162.759	931.71	None
12	3	600	31.76	31.76	1048576	66.34	251.394	4960.547	442.384	None
	4	600	70.34	70.34	1048576	66.57	112.955	4301.322	71.854	None
	5	600	30.9	30.9	1048576	66.33	257.98	4121.884	434.851	None
	6	600	4.11	41.07	10485760	66.71	1937.735	5619.304	904.513	None

Figure 4 – Vdbench results for test cases 7 to 12

The following two figures show the Enterprise Device Manager (EDM) port statistics. Figure 4 shows the ingress and egress port statistics, and Figure 5 shows that there were no errors.

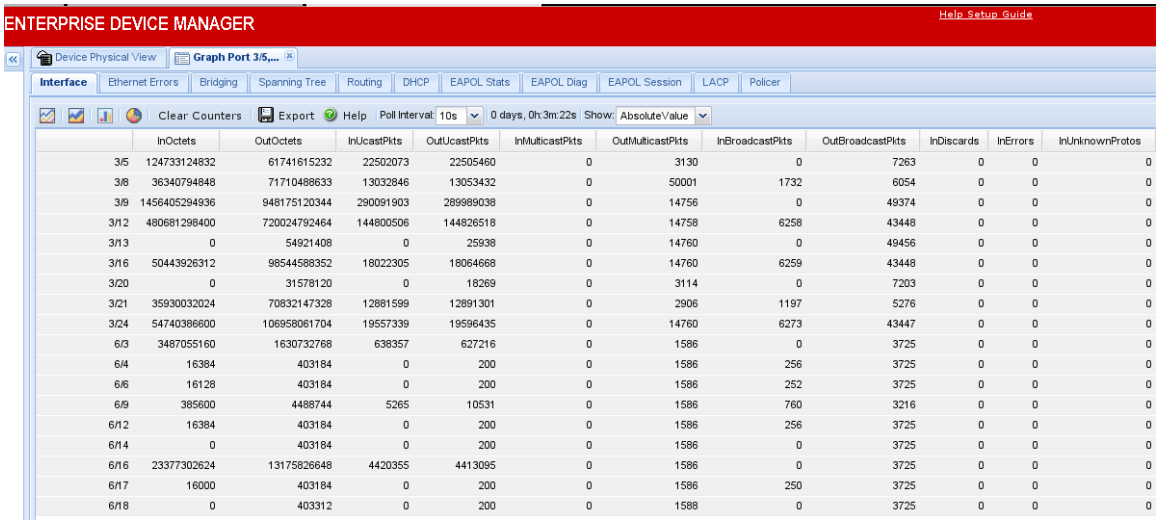

	InOctets	OutOctets	InUcastPkts	OutUcastPkts	InMulticastPkts	OutMulticastPkts	InBroadcastPkts	OutBroadcastPkts	InDiscards	InErrors	InUnknownProtos
3/5	124733124832	61741615232	22502073	22505460	0	3130	0	7263	0	0	0
3/8	36340794848	71710488633	13032846	13053432	0	50001	1732	6054	0	0	0
3/9	1456405294936	948175120344	290091903	289989038	0	14756	0	49374	0	0	0
3/12	480681298400	720024792464	144800506	144826518	0	14758	6258	43448	0	0	0
3/13	0	54921408	0	25938	0	14760	0	49456	0	0	0
3/16	50443926312	98544588352	18022305	18064668	0	14760	6259	43448	0	0	0
3/20	0	31578120	0	18269	0	3114	0	7203	0	0	0
3/21	35930032024	70832147328	12881599	12891301	0	2906	1197	5276	0	0	0
3/24	54740386600	106958061704	19557339	19596435	0	14760	6273	43447	0	0	0
6/3	3487055160	1630732768	638357	627216	0	1586	0	3725	0	0	0
6/4	16384	403184	0	200	0	1586	256	3725	0	0	0
6/6	16128	403184	0	200	0	1586	252	3725	0	0	0
6/9	385600	4488744	5265	10531	0	1586	760	3216	0	0	0
6/12	16384	403184	0	200	0	1586	256	3725	0	0	0
6/14	0	403184	0	200	0	1586	0	3725	0	0	0
6/16	23377302624	13175826648	4420355	4413095	0	1586	0	3725	0	0	0
6/17	16000	403184	0	200	0	1586	250	3725	0	0	0
6/18	0	403312	0	200	0	1588	0	3725	0	0	0

Figure 5 – EDM port statistics (screen 1)

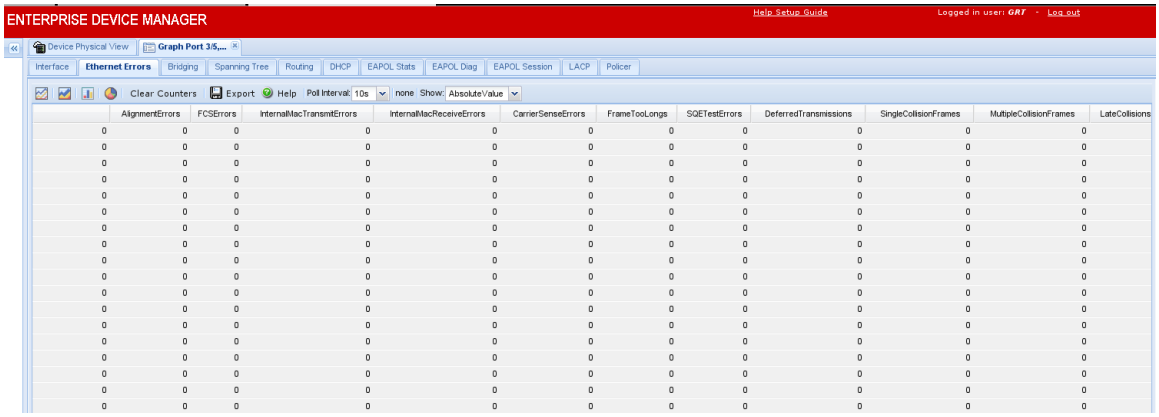

Figure 6 – EDM port statistics (screen 2)

7. Conclusion

The Avaya Data Solutions Test Lab conducted comprehensive tests using a methodology that verified throughput under normal conditions and under various fault conditions. In general the average data rate per PC was approximately 480Mb/s, and the average data transfer size was 1MB. There were no errors observed during the testing.

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.

8. 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 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/.