NetApp
Go further, faster

Technical Report

Using NetApp Unified Connect to Create a Converged Data Center

Freddy Grahn, Chris Lemmons, NetApp
November 2010 | TR-3875

EXECUTIVE SUMMARY

NetApp extends its leadership in Ethernet storage to deliver increased storage flexibility and efficiency with the introduction of NetApp Unified Connect. Unified Connect leverages the NetApp Unified Storage Architecture to support transmission of NAS, iSCSI, and FCoE storage traffic simultaneously over shared 10-Gigabit Ethernet (10GbE) ports using the NetApp unified target adapter. Unified Connect significantly enhances the value proposition of the converged network by offering end-to-end network convergence, from host to storage, in conjunction with FCoE–enabled host adapters and 10GbE switches.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 INTRODUCTION

Today's data centers run multiple parallel networks to accommodate both data and storage traffic. To support these different networks in the data center, administrators deploy separate network infrastructures, including different types of host adapters, connectors and cables, and fabric switches. The use of separate infrastructures increases both capital and operational costs for IT executives. The deployment of a parallel storage network, for example, adds to the overall capital expense in the data center, while the incremental hardware components require additional power and cooling, management, and rack space that negatively affect operational expense.

Additionally, maintaining multiple parallel networks for data and storage traffic can add significantly to management complexity and expense by requiring multiple support groups to maintain the different networks in the data center.

Consolidating these networks in the data center into a unified, integrated infrastructure, referred to as network convergence, reduces both the overall capital expenditure required for network deployment and the operational expenditure for maintaining the infrastructure. This consolidation can also reduce the expense and complexity of managing the data center.

NetApp FAS and V-Series storage systems use a unified storage architecture to simultaneously support all major storage protocols—Network File System (NFS), Common Internet File System protocol (CIFS), Internet Small Computer System Interface (iSCSI), Fibre Channel (FC), and Fibre Channel over Ethernet (FCoE)—in the same storage system. Our unified storage solutions offer you the benefits of simplified data management, investment protection, and reduced total cost of ownership.

NetApp extends its leadership in Ethernet storage to deliver increased storage flexibility and efficiency with the introduction of NetApp Unified Connect. Unified Connect leverages the NetApp Unified Storage Architecture to support transmission of NFS, CIFS, iSCSI, and FCoE storage traffic simultaneously over shared 10-Gigabit Ethernet ports using the NetApp unified target adapter. Unified Connect significantly enhances the value proposition of the converged network by offering end-to-end network convergence, from host to storage, in conjunction with FCoE–enabled host adapters and 10GbE switches.

Unified Connect offers the following benefits:

- True end-to-end network convergence, including IP and FCoE storage traffic
- Reduced overall complexity by reducing the number of management points at the target
- Improved efficiency by increasing the overall bandwidth utilization, allowing a more effective use of your data center infrastructure
- Reduced cabling requirements and significantly improved port utilization
- Agile and high-performing network for increased flexibility and lower capex
- Reduced capex and opex by minimizing the number of switches to purchase and maintain

The converged data center, described in this report, demonstrates how to achieve these benefits for both your existing applications and your future deployments using Unified Connect along with the NetApp unified target adapter and unified storage system. The configuration details and the validation results of the converged data center are also covered in this report.

2 BENEFITS OF A CONVERGED DATA CENTER

To showcase the benefits of Unified Connect, this section compares an existing data center that has grown over the years to a converged data center that uses Unified Connect to reduce its overall complexity. The existing data center, shown in Figure 1, includes multiple applications running different protocols and different virtualization environments across multiple data fabrics. The converged data center, shown in Figure 2, transforms the existing data center into a converged data center consisting of unified target adapters and NetApp unified storage controllers that together support all of the existing network protocols and use fewer wires, switches, and network adapter cards.

2.1 EXISTING DATA CENTER

Figure 1 shows the initial environment. For a number of years, this data center hosted a production Oracle® Database environment using 4Gb FC connectivity. Recently, applications were added and configured that run in both VMware® vSphere™ 4.1 and Microsoft® Hyper-V™ R2 virtualized environments.

Figure 1) Existing data center diagram.

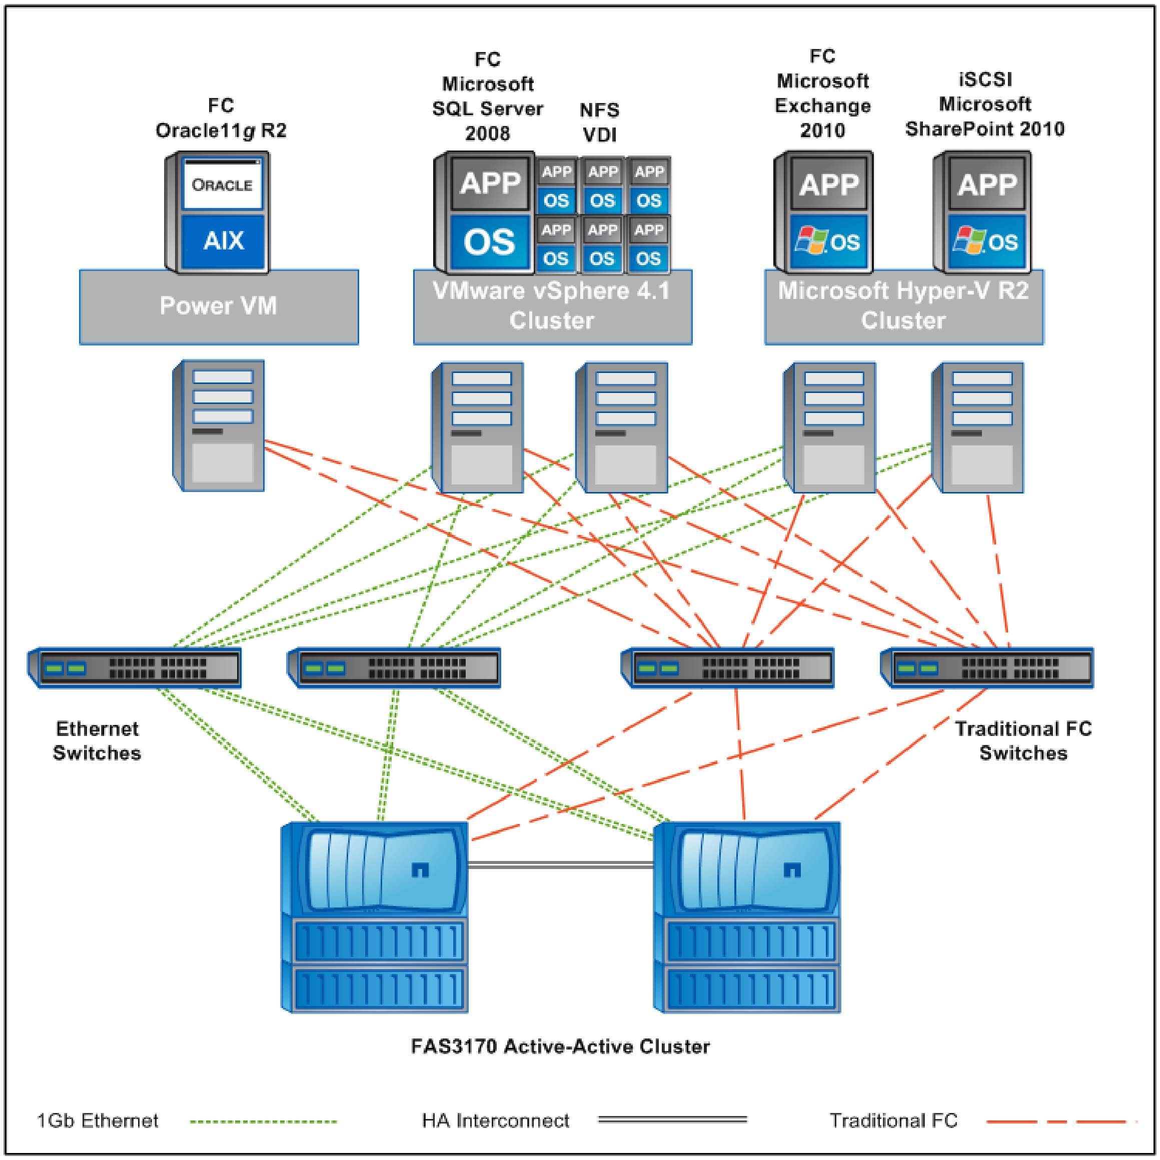

The following applications are configured in the existing vSphere 4.1 environment:

- Microsoft SQL Server® 2008 connected to the storage using standard 4Gb FC
- 50-seat VDI environment containing 50 Windows® 7 desktops connected to the storage using NFS v3 across a single Gigabit Ethernet connection

The following applications are configured in the existing Microsoft Hyper-V R2 environment:

- Microsoft Exchange 2010 with 1,000 heavy users and 1GB mailboxes using standard 4Gb FC
- Microsoft Office SharePoint® 2010 using iSCSI across a single Gigabit Ethernet connection

The four servers configured in the vSphere and Hyper-V R2 environments are identical HP ProLiant DL360 G6 servers. Each server is configured with 32GB of RAM and 2 x quad-core Xeon™ 2.133Ghz processors. For connectivity, each of these servers contains a dual-ported 4Gb FC adapter and a dual-ported Gigabit Ethernet adapter.

A NetApp FAS3170A unified storage system running Data ONTAP® 7.3.2 provides storage support for this data center. Each of the NetApp FAS3170A controllers is configured with built-in 4Gb FC adapters and multiple Gigabit Ethernet ports. The controllers are connected in a standard failover configuration to the data center servers using standard 4Gb FC and Gigabit Ethernet switches.

As the data center expands, more switches and server adapters are required to support the configuration. In addition, the data center potentially requires different network administrators to support the FC and Gigabit Ethernet infrastructures.

2.2 CONVERGED DATA CENTER

Figure 2 depicts how we transformed the existing data center into a modern converged data center that greatly simplifies how the older Oracle and newer virtualized environments are connected and managed.

We replaced the 4Gb FC and Gigabit Ethernet adapters in the four servers hosting the VMware vSphere 4.1 and Hyper-V R2 environments with a single dual-port QLogic® 8152 10Gb Converged Network Adapter (CNA). The QLogic 8152 card allows the vSphere and Hyper-V R2 servers to transmit normal TCP/IP and FC traffic simultaneously across the same high-performance 10GbE wire. The QLogic 8152 provides a number of benefits including, but not limited to, cost savings through reduced adapter, switch, cabling, power, cooling, and management requirements.

Additionally, we replaced the 4Gb FC adapter in the IBM® p520 server with a newer 8Gb FC adapter to improve the performance of the production Oracle Database in the new converged data center. This allowed us to preserve the existing database instance that is currently installed in the data center running in a hardware virtualized AIX environment. By connecting this existing Oracle configuration to the new virtualized configurations using the same converged networking infrastructure, we demonstrate that customers don't have to abandon their significant investments in storage infrastructure and applications to build for the future using more modern converged networks and NetApp unified storage systems.

To allow the NetApp FAS3170A unified storage systems to support this new converged environment, we upgraded the storage to run Data ONTAP 8.0.1. As previously mentioned, NetApp unified storage supports CIFS, NFS, iSCSI, FC, and FCoE using a single platform. This makes NetApp unified storage systems an ideal option for use in this type of converged environment. We then configured each of the NetApp FAS3170A controllers with a single dual-ported unified target adapter. The unified target adapter allows the NetApp unified storage system to process all of the different types of traffic generated by the converged data center (FCoE, NFS, iSCSI, and FC traffic initiated by the 8Gb FC adapters installed in the IBM p520) using a single card and wire installed in the storage controllers.

All of these server and storage systems are then connected using a pair of Cisco® Nexus 5010 switches. The Cisco Nexus 5000 Series switches provide innovative support for a standards-based, multilayer, multiprotocol, and multipurpose Ethernet-based fabric as described in the converged data center.

To connect the Cisco switches to the storage and servers, we used Amphenol SFP+ cable assemblies. Amphenol's SFP+ cable assemblies are high performance, cost effective I/O solutions for 10Gb Ethernet and 8GB Fibre Channel applications. SFP+ direct-attached copper cables allow hardware manufacturers to achieve high port density, configurability, and utilization at a very low cost and reduced power budget. Amphenol's high speed cable assemblies meet and exceed Gigabit Ethernet and Fibre Channel industry standard requirements for performance and reliability.

Figure 2) Converged data center diagram.

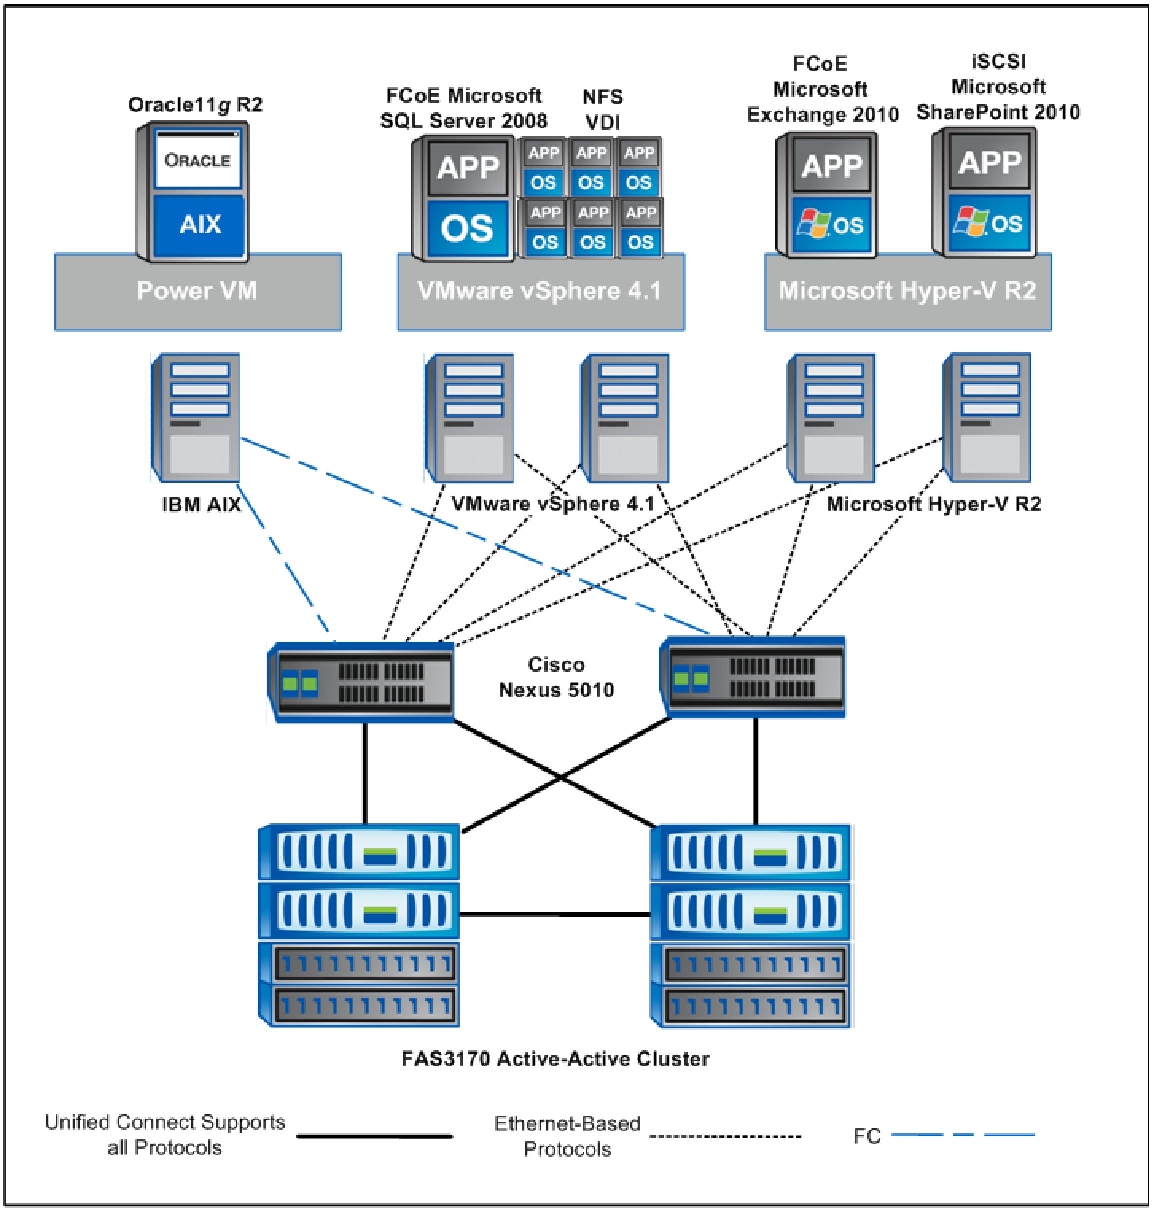

As shown in Figure 2, moving to this converged data center enabled us to:

- Replace 16 ports (4 x dual port Gigabit Ethernet and 4 x dual port 4Gb FC adapters) with 8 ports (4 x dual port unified target adapters).
- Significantly improve the overall bandwidth capacity for FC, FCoE, iSCSI, and NFS by replacing the 4Gb FC adapter in the AIX server with an 8Gb FC adapter and replacing the 4Gb FC and Gigabit Ethernet adapters in the virtualization servers with one 10Gb unified target adapter.
- Cut the required number of switches in half by using the Cisco Nexus 5010 product.
- Provide a common, unified fabric for use by all of the data center applications and the NetApp unified storage controllers.

2.3 PERFORMANCE TESTS

After configuring the converged data center, we validated the overall configuration by conducting a series of performance tests using a set of application-specific performance measurement tools. The performance tests are not designed to measure the maximum performance capabilities of the servers or storage. They are conducted to demonstrate that NetApp unified storage configured with Unified Connect is capable of providing acceptable performance in a multiapplication environment in which multiple protocols are used.

Performance tests were conducted both individually for the Oracle, SQL Server, SharePoint, Exchange, and virtual desktop infrastructure (VDI), and also by running all these applications at the same time. Regardless whether the tests were running individually or simultaneously using multiple protocols, our results indicated that NetApp unified storage provided great overall performance while maintaining acceptable latencies that were well within the tolerances of each of the applications.

The following sections provide a detailed discussion on how we configured each of the applications deployed in the converged data center and the application-specific tools used for the performance validation testing.

3 DATA CENTER CONFIGURATION DETAILS

This section describes the three areas that make up the converged data center, including applications running in both vSphere 4.1 and Hyper-V R2 environments as well as the Oracle11*g* R2 database environment.

3.1 CISCO NEXUS CONFIGURATION DETAILS

In this solution, the network is composed of two Cisco Nexus 5010 switches for managing both the FC and IP-based network traffic generated in the converged data center. We performed the following operations on the Cisco Nexus network:

- Configured FC ports
- Configured a 10Gb connection between the two Cisco Nexus 5010 switches
- Enabled a Virtual Port Channel (vPC) between the two Cisco Nexus 5010 switches
- Licensed FCoE on both Nexus switches

Cisco Nexus 5010 switches support both 8Gb FC and 10Gb IP. They also support 1Gb modules allowing other Cisco switches to be used in conjunction with the Cisco Nexus 5010s. This support enables further scale-out of the virtualization and storage network.

3.2 ORACLE DATABASE CONFIGURATION DETAILS

The standalone Oracle database represents an existing production configuration with a significant investment of resources for software licenses and programming. Therefore, this existing architecture was not abandoned as a result of the change to the unified fabric. Our demonstration shows how this configuration can be connected into the new network fabric to enable the future expansion of the data center while continuing to leverage the existing configurations.

In this case, Oracle11*g* R2 is deployed on an IBM p520 server running AIX 6.1 TL5 SP2. This server contains a standard dual-ported 8Gb/s FC adapter. The Oracle database software is installed in a single logical partition (LPAR). Two virtual FC adapters are configured on the LPAR.

As shown in Figure 2, the FC ports on the p520 are connected to the Cisco 5010 switches using the standard FC ports on the switch.

STORAGE ARCHITECTURE FOR ORACLE

Figure 3 shows the NetApp storage aggregate and volume layout for hosting the different data components for the AIX configuration running Oracle. NetApp aggregates provide a large virtualized pool of storage capacity and disk input/output operations per second (IOPS) to be used on demand by all the applications hosted in the aggregate.

Figure 3) AIX volume layout.

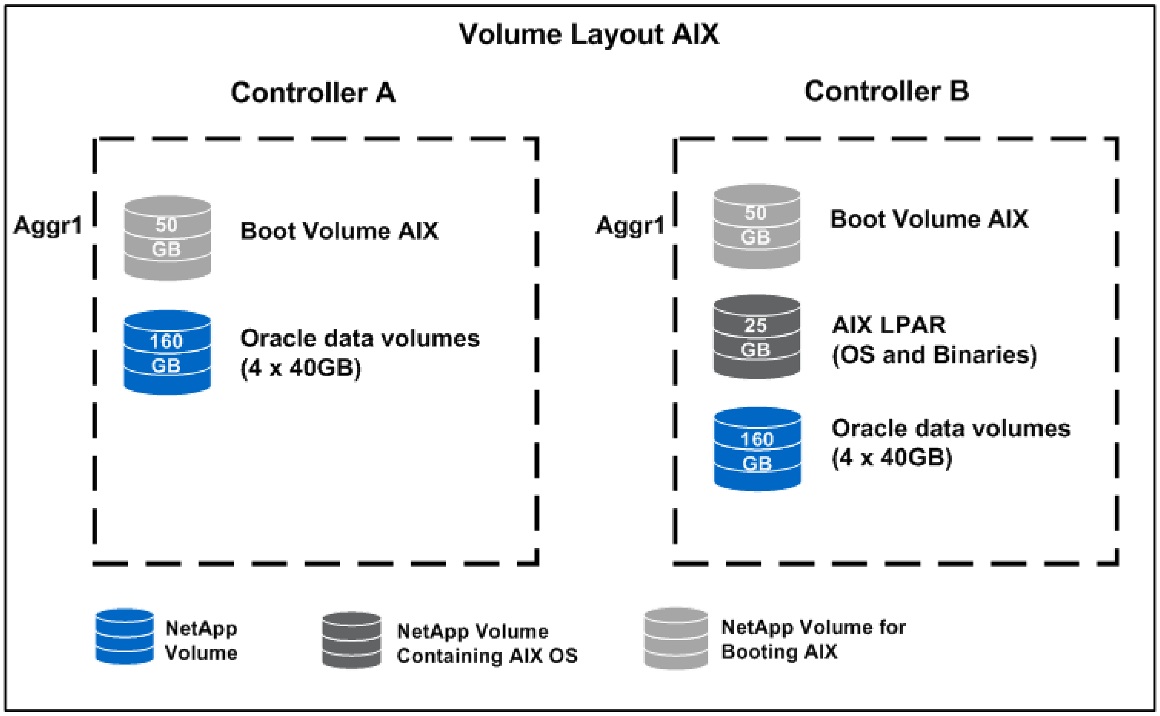

The aggregate sizing is based on the storage requirements for all the applications needed to meet the storage capacity requirement of an assumed workload. When sizing for your environment, consult with your NetApp systems engineer (SE) about the exact storage configuration based on your individual requirements.

Note: In this solution, a single aggregate is created on each of the storage systems. Each aggregate contains 23 data disks with 1 spare disk. All of the volumes are thin provisioned to use the capacity on demand.

We carefully separated data files and log files onto different aggregates. The boot logical unit number (LUN) for the LPAR is created in both aggregates and acts as a failover boot LUN. The binaries for Oracle and AIX 6.1 are also stored on a NetApp LUN.

ORACLE APPLICATION CONFIGURATION

Oracle11*g* R2

Figure 4 shows the layout for the different data components of the Oracle11*g* R2 database. The Oracle database and log files are hosted on FC LUNs on separate volumes and aggregates. The Oracle binaries and AIX boot code are contained in their own LUNs as well.

Figure 4) Oracle application layout.

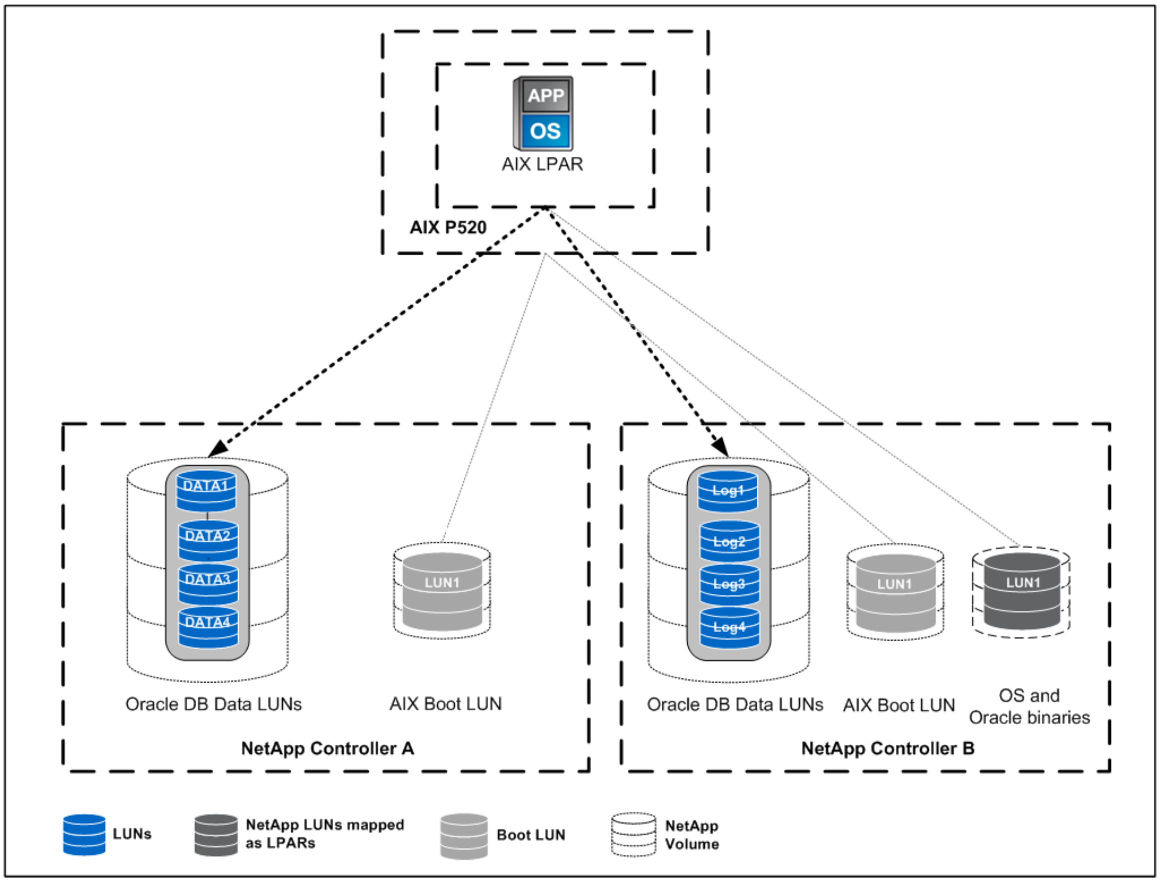

3.3 ESX 4.1 CONFIGURATION DETAILS

This section provides details on how vSphere and the associated hosted applications are configured and includes both the vSphere networking and NetApp storage details.

VIRTUAL MACHINES USED

Application Virtual Machines

We used a single virtual machine (VM) to configure Microsoft SQL Server 2008 under vSphere. For the VDI configuration, we used a total of 50 VMs. Both configurations simulate a real-world customer environment. The details of the SQL Server 2008 and VDI configurations are as follows:

- SQL Server 2008 using FCoE
  - 1 SQL Server VM (Win2k8R2, 4vCPU, 4 GB); 1,500 users across 3 databases (DBs)
  - Storage: User DB – 2.5TB, Logs – 350GB; Temp DB – 300GB; System DB – 300GB
- VDI using NFS
  - 50 VMs – 50 Windows 7 clients (1 vCPU, 1,024MB RAM each)
  - Storage: 2 NFS datastores of 500GB each

For the SQL Server 2008 VM using FCoE, we applied all of the NetApp recommended Fibre Channel Protocol (FCP) settings (refer to TR-3749: NetApp and VMware vSphere Storage Best Practices) by using the NetApp VSC vCenter™ plug-in, directly from the vCenter GUI.

Utility Virtual Machines

Multiple load generators and test tools were used to validate the setup. SQLIO was loaded onto an ESX VM to generate a load similar to that of SQL Server 2008 during the performance validation phase. For VDI testing, three VMs were used to install the VMware View Connection Broker, VMware View Client, and VMware vCenter 4.1. The RAWC tool for VDI was used to generate the load from all 50 individual desktops. All utility VMs were provisioned and stored on a volume that was configured to use FCoE.

- Utility VM accessed using FCoE:
  - VMware vCenter 4.1 (Win2k3R2, required for VDI)
  - VMware View Connection Broker (Win2k3R2, required for VDI)
  - SQLIO VM

ESX NETWORK ARCHITECTURE

Figure 5 shows the virtual network layout for each ESX host. Each ESX host has two 10GbE ports configured into different port groups. We applied all of the other FC configuration best practices highlighted in [TR-3749: NetApp and VMware vSphere Storage Best Practices](#).

Figure 5) ESX host network architecture.

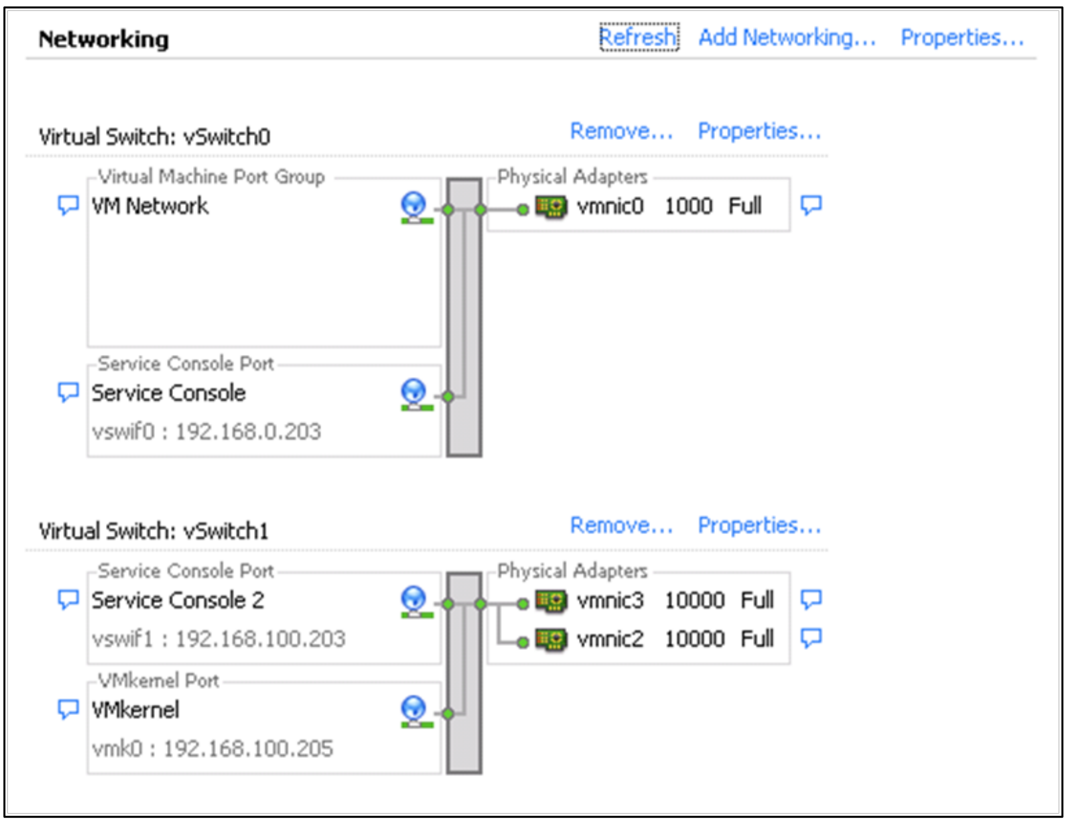

STORAGE ARCHITECTURE FOR ESX 4.1 APPLICATIONS

Figure 6 shows the NetApp storage aggregate and volume layout for hosting the different data components for every VM running in the vSphere environment. NetApp aggregates provide a large virtualized pool of storage capacity and disk IOPS to be used on demand by all the VMs hosted in the aggregate. NetApp aggregates can be compared to VMware virtualization in which CPU and memory resources are pooled and leveraged on demand.

Figure 6) ESX storage layout.

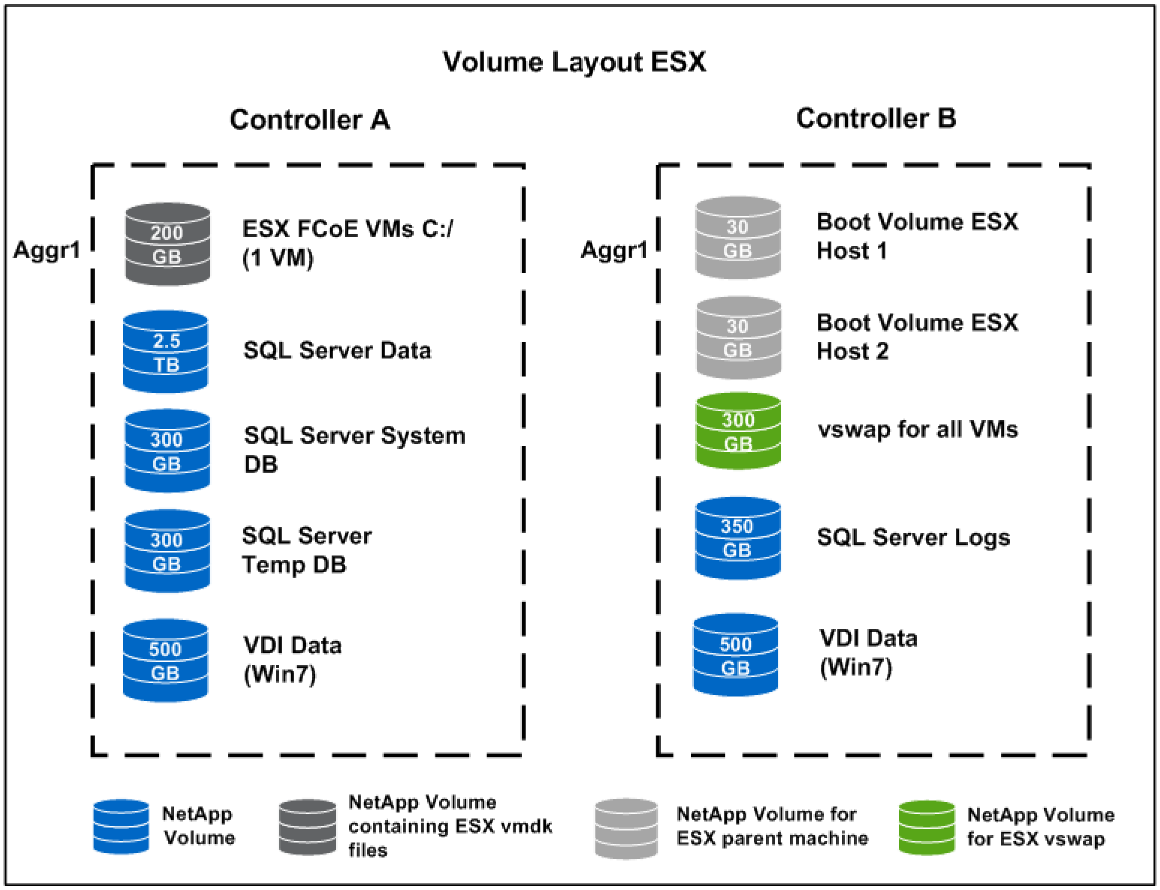

The aggregate sizing is based on the storage requirements for all the applications to meet the storage capacity requirement of an assumed workload. When sizing for your environment, consult with your NetApp SE about the exact storage configuration based on your individual requirements.

Note: In this solution, a single aggregate is created on each of the storage systems. Each aggregate contains 23 data disks with 1 spare disk. All of the volumes are thin provisioned to use the capacity on demand.

The data files and log files were carefully separated onto different aggregates. We also separated the VM C:\ drives from the VM vswap files. The Windows 7 VDI VMs were also spread across different aggregates on different storage controllers.

Each VM has a 24GB C: drive (minimum requirements for Windows Server 2008 per the VMware Guest Operating System Installation Guide) with the vmdk hosted on the VMFS datastore. The application server (SQL Server) database and log drives are hosted on FC-based RDM LUNs. These LUNs are directly created and connected inside the guest VMs using NetApp SnapDrive® 6.3 software. This provides the flexibility to leverage the NetApp and Microsoft application-integrated SnapDrive products to achieve granular LUN management.

APPLICATION CONFIGURATION IN ESX 4.1

Microsoft SQL Server 2008 Layout

Figure 7 shows the layout for the different data components of Microsoft SQL Server 2008. The SQL database and log files are hosted on FC RDM LUNs on separate volumes and aggregates. The binaries, boot LUN, and VM vswap are also contained in their own LUNs.

Figure 7) Microsoft SQL Server layout.

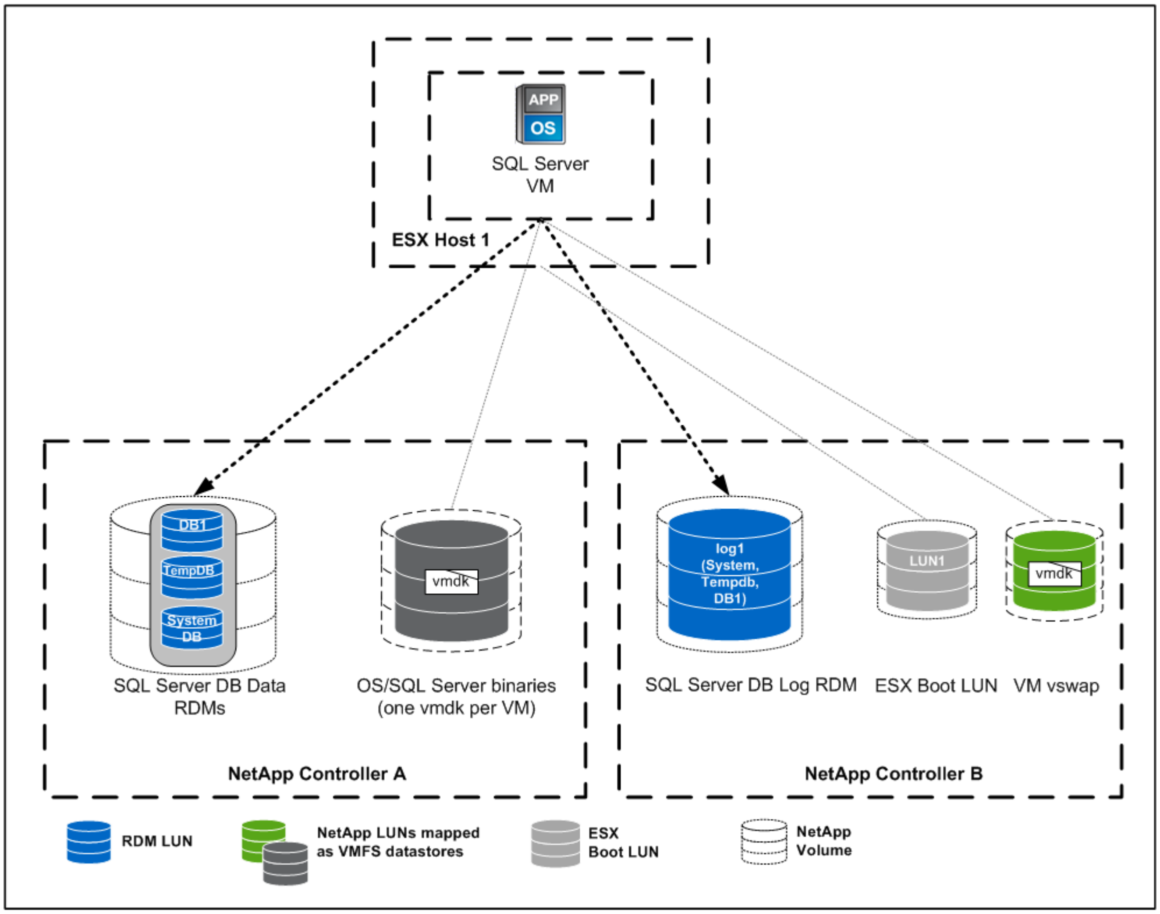

VDI Layout

Figure 8 shows the layout for the different data components of the VDI configuration. We used a total of 50 VMs to create this environment (50 Windows 7 clients). Each of the clients stored the VM vswap in a single separate LUN. Two separate FlexVol® volumes were used to store the VDI VMs, and a separate boot LUN was used to boot ESX. The VDI VMs were stored on a NFS v3 datastore.

Figure 8) VDI layout.

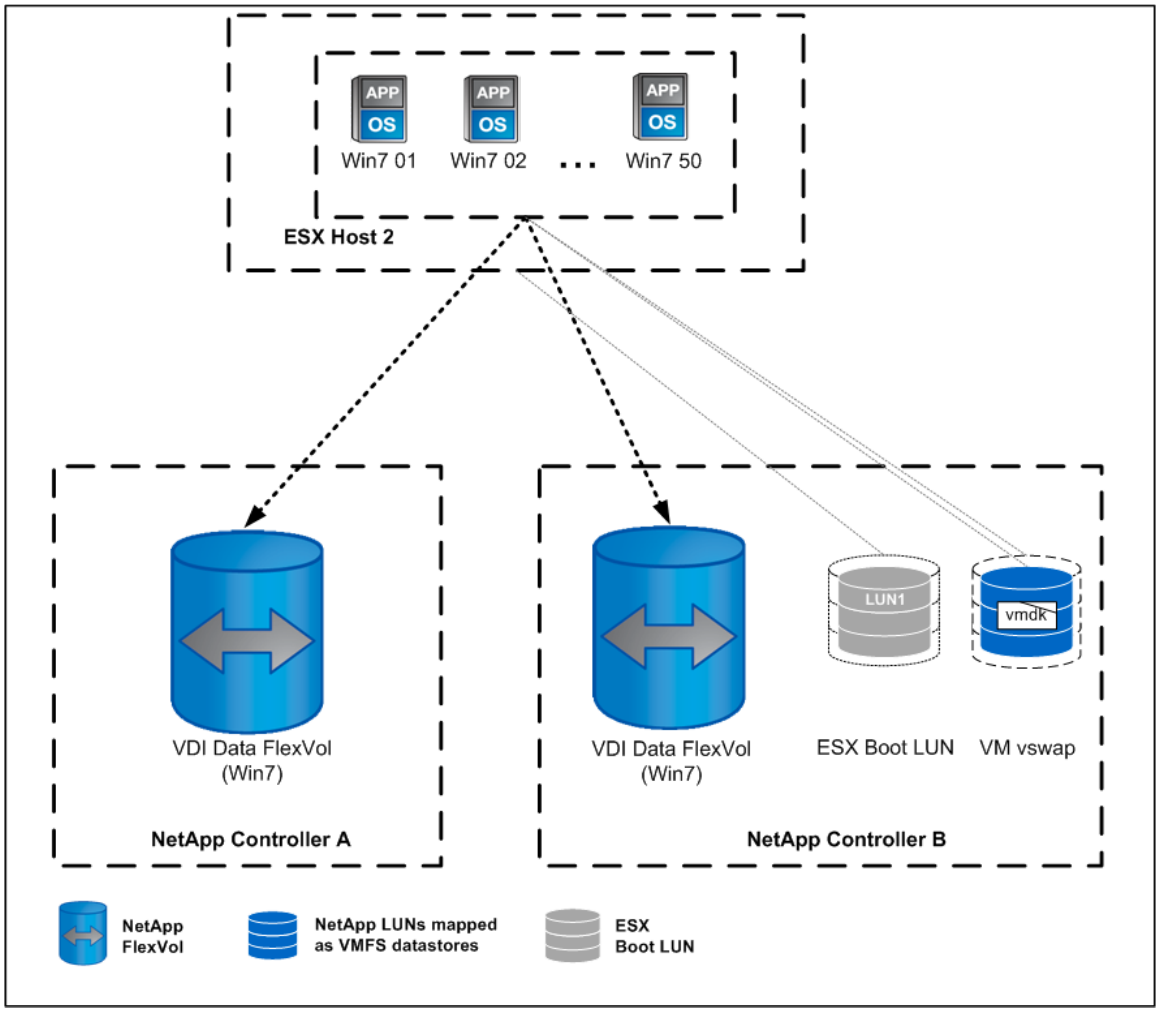

3.4 HYPER-V R2 CONFIGURATION DETAILS

This section provides the configuration details for hosting the Hyper-V R2 VMs and the associated applications. Details on both the Hyper-V R2 networking and NetApp storage are also included.

VIRTUAL MACHINES USED

Application Virtual Machines

The solution described in this section uses a total of six VMs. This configuration simulates a real-world environment with supporting utility VMs and includes the primary Microsoft application servers.

- Microsoft Exchange 2010 using FCoE
  - Mailbox Server (Win2k8R2, 4vCPU, 20GB RAM) – 1,000 users at 1GB/user
  - 1 Hub transport and Client Access Server (Win2k8R2, 2vCPU, 4GB RAM)
  - Storage: DB – 2.0TB, Logs – 150GB (rotatable)
- Microsoft Office SharePoint Server 2010 using iSCSI
  - 1 Web Server for SharePoint installation (Win2k8R2, 2vCPU, 4GB RAM) of 3,000 users with 25% active

– 1 SQL Server 2008 and Index Server (Win2k8R2, 4vCPU, 4GB RAM)
– Storage: System DB – 300GB; Temp DB – 300GB, Logs – 350GB; SharePoint DB – 500GB

Utility Virtual Machines

We used multiple load generators and test tools to validate the setup. To test Microsoft Exchange 2010, we loaded the Microsoft LoadGen tools onto a separate Hyper-V R2 VM. For SharePoint, we used one VM to install Avepoint SharePoint Test Environment Creator and Usage Simulator. All utility VMs were provisioned and stored on a volume that was configured to use FCoE.

- Utility VMs
  - 1 LoadGen VM (Win2k8R2, 2vCPU, 2GB)
  - 1 VM (Win2k8R2, 2vCPU, 2GB) for Avepoint SharePoint Test Environment Creator and Usage Simulator

STORAGE ARCHITECTURE FOR HYPER-V R2 APPLICATIONS

Figure 9 shows the two NetApp storage aggregates (one aggregate per controller) and volume layout for hosting the different data components for all of the VMs in the Hyper-V R2 environment. NetApp aggregates provide a large virtualized pool of storage capacity and disk IOPS to be used on demand by all the VMs hosted in the aggregate. This can be compared to the Hyper-V virtualization in which CPU and memory resources are pooled and leveraged on demand.

Figure 9) Hyper-V storage layout.

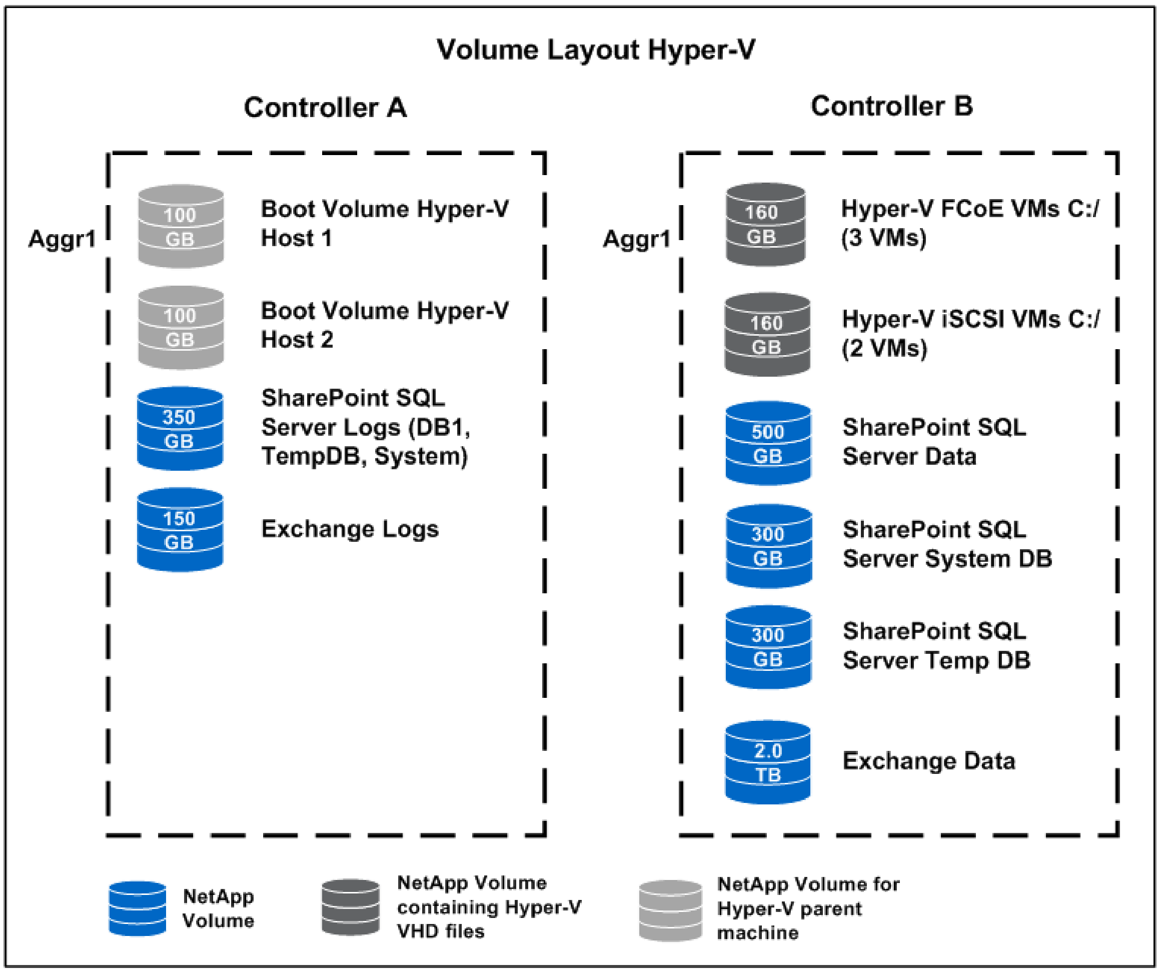

The aggregate sizing is based on the storage requirements for all the applications to meet the storage capacity requirement of an assumed workload. When sizing for your environment, consult with your NetApp SE about the exact storage configuration based on your individual requirements.

Note: In this solution, a single aggregate is created on each of the storage systems. Each aggregate contains 23 data disks with 1 spare disk. All of the volumes are thin provisioned to use the capacity on demand.

Data files and log files are carefully separated onto different aggregates.

Each VM has a 24GB C: drive (maintaining the same drive space as ESX) with each vhd hosted on a NetApp LUN. All application servers (Exchange, SharePoint, SQL Server) and log drives are hosted on FC–based pass through LUNs, directly created and connected inside the guest VMs using NetApp SnapDrive 6.3 software. This provides the flexibility to leverage the NetApp and Microsoft application-integrated SnapDrive products to achieve granular LUN management.

APPLICATION CONFIGURATION IN HYPER-V R2

Microsoft Exchange 2010 Layout

Figure 10 shows the layout for the different data components of Microsoft Exchange 2010. The Exchange log and data files are stored on Hyper-V pass through LUNs within separate volumes and aggregates. The binaries and boot LUN are also contained in their own separate LUN.

Figure 10) Microsoft Exchange layout.

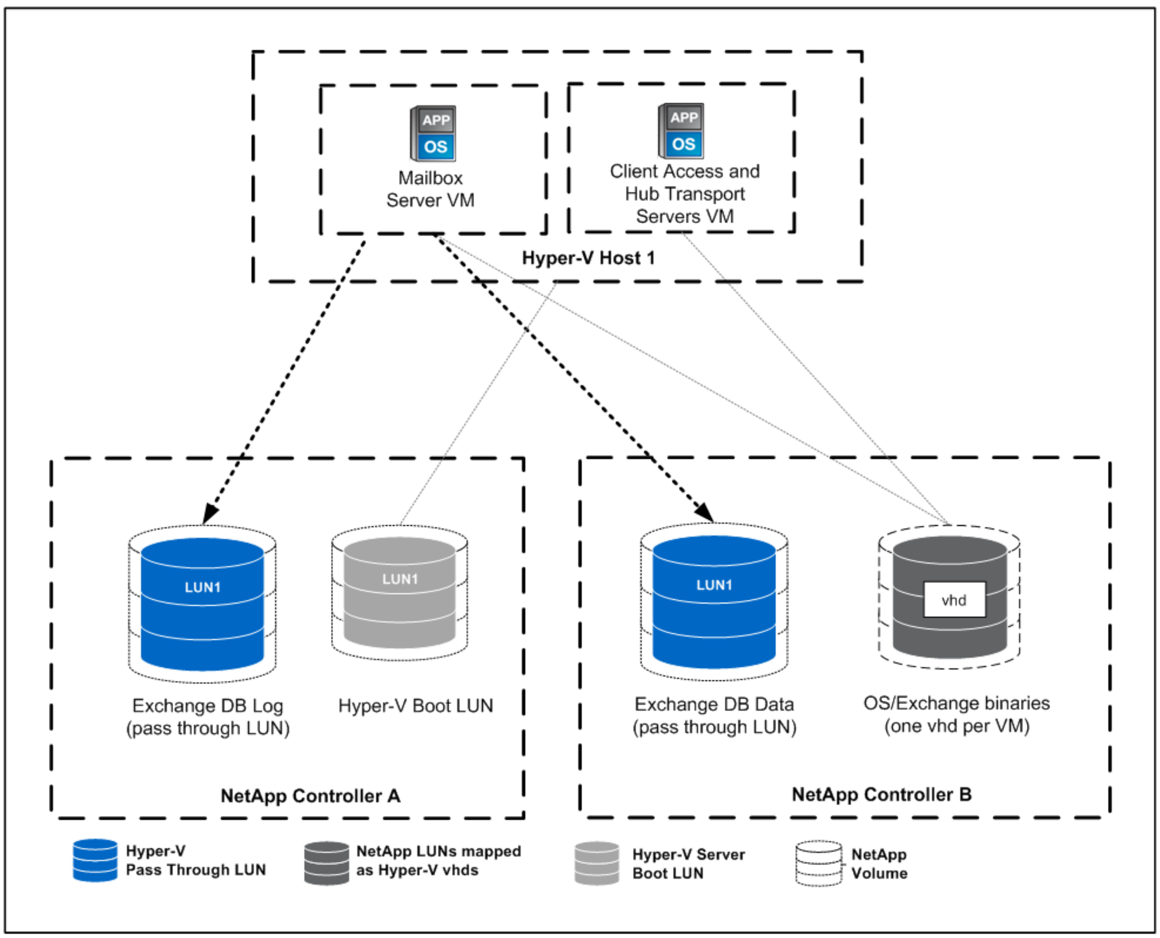

Microsoft SharePoint 2010 Layout

Figure 11 shows the layout for the different data components of Microsoft SharePoint 2010. A VM is set up for SharePoint's Web front end, and a separate VM is set up for the SQL Server 2008 database. The database files and log files are hosted on separate LUNs and on separate volumes within separate aggregates. In addition, the Master DB and Temp DB each have their own LUN and volume separate from the application database. The log files are all placed into the same LUN and volume.

Figure 11) Microsoft SharePoint 2010 layout.

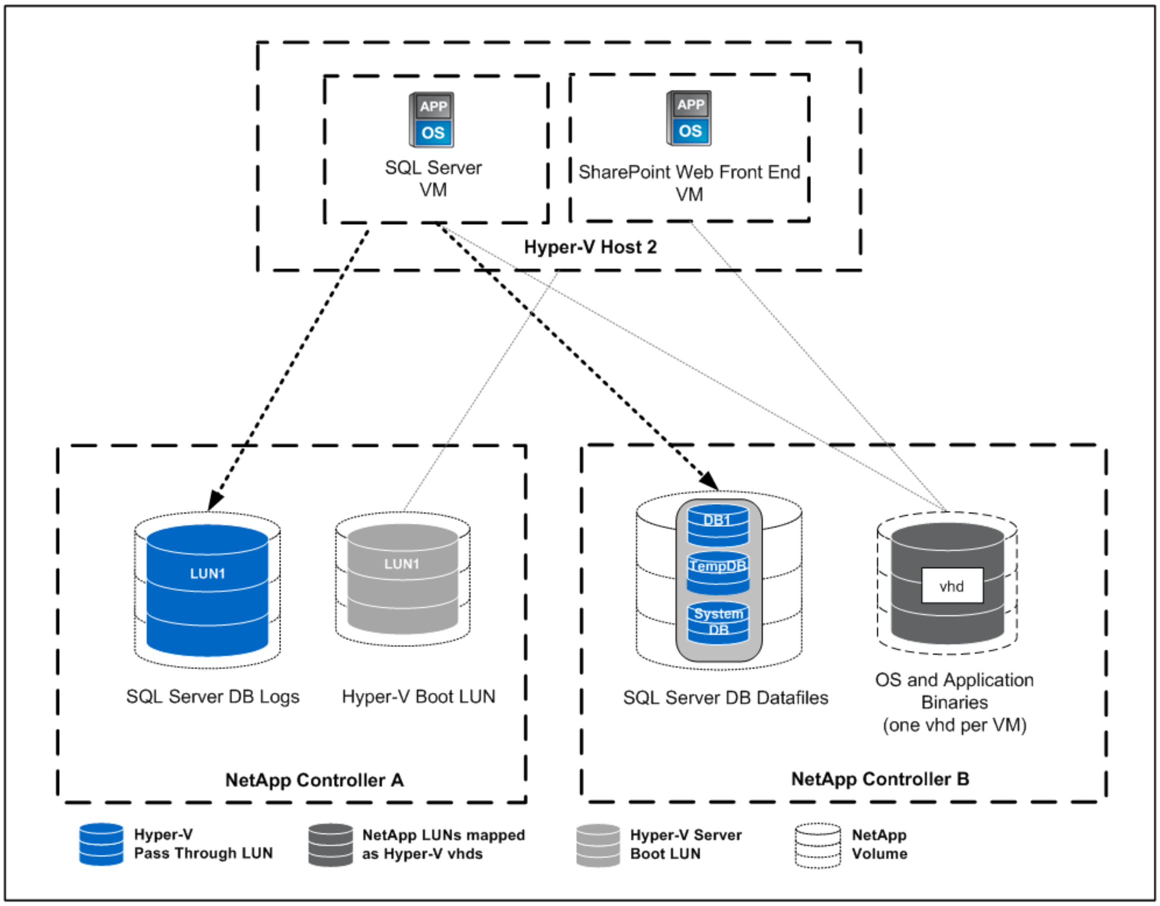

4 VALIDATING THE CONVERGED DATA CENTER

After configuring the converged data center (as described in the previous section), we validated the overall configuration by conducting a series of performance tests using a set of application-specific performance measurement tools. The performance tests were not designed to measure the maximum performance capabilities of the servers or storage. They were conducted to demonstrate that NetApp unified storage, configured with Unified Connect using the unified target adapter, is capable of providing an acceptable level of performance in a multiapplication environment in which multiple protocols are used.

We conducted individual performance tests for Oracle, SQL Server, SharePoint, Exchange, and VDI. We also conducted the tests on all of these applications simultaneously. For both sets of tests, we found that NetApp unified storage provided great overall performance while maintaining acceptable latencies well within the recommended tolerances of each of the applications under test.

We used the following application-specific load-generation tools to validate the performance:

- **Oracle11*g* R2.** The Orion load generator was used to simulate a typical Oracle environment. The workload generator simulated a typical Decision Support Systems (DSS) environment of 1,000 active users.
- **VDI.** The VMware Reference Architecture Workload Simulator (RAWC) tool was used to generate load from all VDI desktops. The user workloads that we simulated used standard Microsoft Office 2007 applications, Microsoft Internet Explorer®, and Windows Media Player in a 50-user environment.
- **Microsoft Exchange 2010.** The Microsoft Exchange Load Generation Tool (LoadGen) was used to simulate a 1,000-user mail profile with 1GB per mailbox.
- **Microsoft SQL Server 2008.** The Microsoft SQLIOSim utility was used to mimic a typical SQL Server workload. The workload generated simulated a typical Online Transaction Processing (OLTP) environment with 1,500 active users.
- **Microsoft SharePoint 2010.** AvePoint SharePoint Test Environment Creator and Usage Simulator tools were used to populate and test the SharePoint environment. The user workload tested was typical of a 1,000-user environment.

Figure 12 shows all protocols accessing the two NetApp FAS3170A unified storage controllers with the generated load when all applications are running simultaneously. The data is from the output of the `sysstat` command that is available on all NetApp unified storage products. Each of the images was captured on one of the FAS3170A unified storage controllers. Together, these images provide a picture of the controllers handling the FC, FCoE, iSCSI, and NFS traffic generated by the converged data center through a single unified target adapter installed in each of the controllers. The NFS, FCP, and iSCSI columns are highlighted in Figure 12. Each of these columns displays the total number of IOPS generated by the applications hosted in the converged data center, for each protocol on each storage controller, at one-second intervals over approximately a 20-second duration.

Figure 12) Screen output of results.

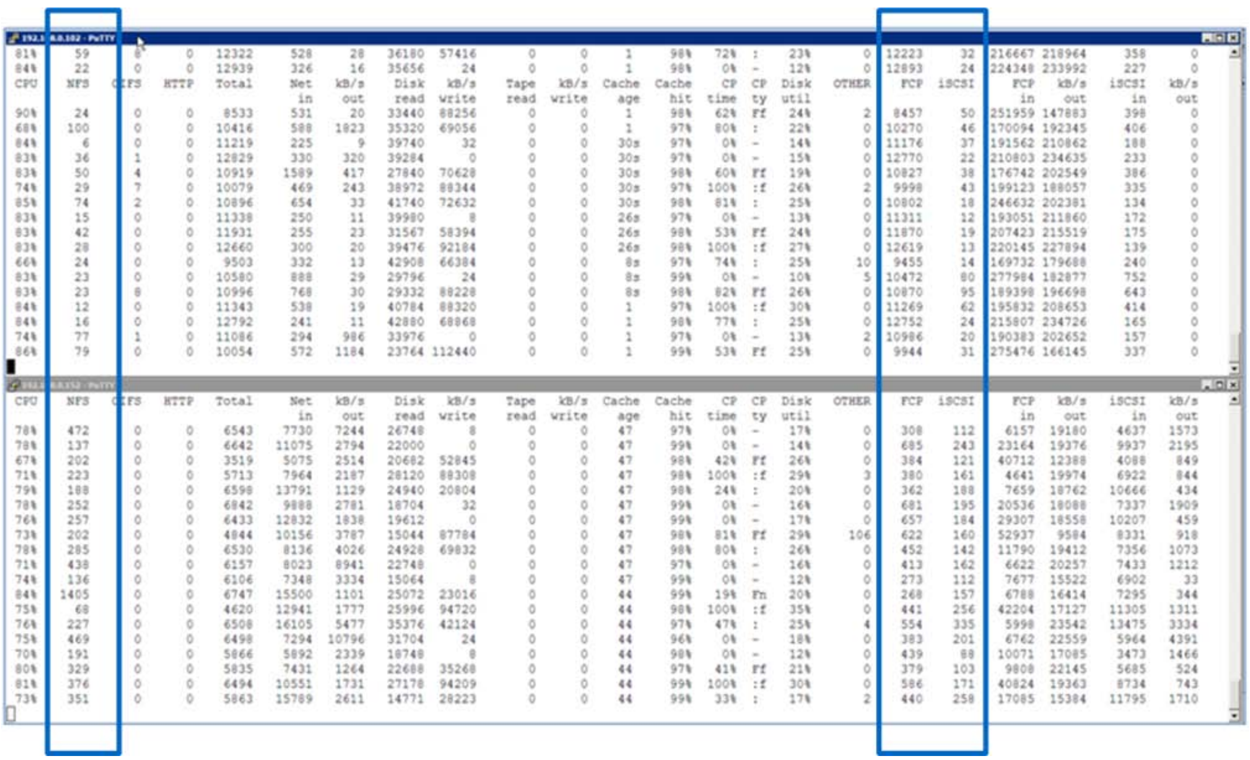

81%	59	8	0	12322	528	28	36180	57416	0	0	1	98%	72%	:	23%	0	12223	32	216667	218964	358	0
84%	22	0	0	12939	326	16	35656	24	0	0	1	98%	0%	-	12%	0	12893	24	224348	233992	227	0
CPU	NFS	CIFS	HTTP	Total	Net	kB/s	Disk	kB/s	Tape	kB/s	Cache	Cache	CP	CP	Disk	OTHER	FCP	iSCSI	FCP	kB/s	iSCSI	kB/s
					in	out	read	write	read	write	age	hit	time	ty	util				in	out	in	out
90%	24	0	0	8533	531	20	33440	88256	0	0	1	98%	62%	Ff	24%	2	8457	50	251959	147883	398	0
68%	100	0	0	10416	588	1823	35320	69056	0	0	1	97%	80%	:	22%	0	10270	46	170094	192345	406	0
84%	6	0	0	11219	225	9	39740	32	0	0	30s	97%	0%	-	14%	0	11176	37	191562	210862	188	0
83%	36	1	0	12829	330	320	39284	0	0	0	30s	97%	0%	-	15%	0	12770	22	210803	234635	233	0
83%	50	4	0	10919	1589	417	27840	70628	0	0	30s	98%	60%	Ff	19%	0	10827	38	176742	202549	386	0
74%	29	7	0	10079	469	243	38972	88344	0	0	30s	97%	100%	:f	26%	2	9998	43	199123	188057	335	0
85%	74	2	0	10896	654	33	41740	72632	0	0	30s	98%	81%	:	25%	0	10802	18	246632	202381	134	0
83%	15	0	0	11338	250	11	39980	8	0	0	26s	97%	0%	-	13%	0	11311	12	193051	211860	172	0
83%	42	0	0	11931	255	23	31567	58394	0	0	26s	98%	53%	Ff	24%	0	11870	19	207423	215519	175	0
83%	28	0	0	12660	300	20	39476	92184	0	0	26s	98%	100%	:f	27%	0	12619	13	220145	227894	139	0
66%	24	0	0	9503	332	13	42908	66384	0	0	8s	97%	74%	:	25%	10	9455	14	169732	179688	240	0
83%	23	0	0	10580	888	29	29796	24	0	0	8s	99%	0%	-	10%	5	10472	80	277984	182877	752	0
83%	23	8	0	10996	768	30	29332	88228	0	0	8s	98%	82%	Ff	26%	0	10870	95	189398	196698	643	0
84%	12	0	0	11343	538	19	40784	88320	0	0	1	97%	100%	:f	30%	0	11269	62	195832	208653	414	0
84%	16	0	0	12792	241	11	42880	68868	0	0	1	98%	77%	:	25%	0	12752	24	215807	234726	165	0
74%	77	1	0	11086	294	986	33976	0	0	0	1	97%	0%	-	13%	2	10986	20	190383	202652	157	0
86%	79	0	0	10054	572	1184	23764	112440	0	0	1	99%	53%	Ff	25%	0	9944	31	275476	166145	337	0

CPU	NFS	CIFS	HTTP	Total	Net	kB/s	Disk	kB/s	Tape	kB/s	Cache	Cache	CP	CP	Disk	OTHER	FCP	iSCSI	FCP	kB/s	iSCSI	kB/s
					in	out	read	write	read	write	age	hit	time	ty	util				in	out	in	out
78%	472	0	0	6543	7730	7244	26748	8	0	0	47	97%	0%	-	17%	0	308	112	6157	19180	4637	1573
78%	137	0	0	6642	11075	2794	22000	0	0	0	47	99%	0%	-	14%	0	685	243	23164	19376	9937	2195
67%	202	0	0	3519	5075	2514	20682	52845	0	0	47	98%	42%	Ff	26%	0	384	121	40712	12388	4088	849
71%	223	0	0	5713	7964	2187	28120	88308	0	0	47	98%	100%	:f	29%	3	380	161	4641	19974	6922	844
79%	188	0	0	6598	13791	1129	24940	20804	0	0	47	98%	24%	:	20%	0	362	188	7659	18762	10666	434
78%	252	0	0	6842	9888	2781	18704	32	0	0	47	99%	0%	-	16%	0	681	195	20536	18088	7337	1909
76%	257	0	0	6433	12832	1838	19612	0	0	0	47	99%	0%	-	17%	0	657	184	29307	18558	10207	459
73%	202	0	0	4844	10156	3787	15044	87784	0	0	47	98%	81%	Ff	29%	106	622	160	52937	9584	8331	918
78%	285	0	0	6530	8136	4026	24928	69832	0	0	47	98%	80%	:	26%	0	452	142	11790	19412	7356	1073
71%	438	0	0	6157	8023	8941	22748	0	0	0	47	97%	0%	-	16%	0	413	162	6622	20257	7433	1212
74%	136	0	0	6106	7348	3334	15064	8	0	0	47	99%	0%	-	12%	0	273	112	7677	15522	6902	33
84%	1405	0	0	6747	15500	1101	25072	23016	0	0	44	99%	19%	Fn	20%	0	268	157	6788	16414	7295	344
75%	68	0	0	4620	12941	1777	25996	94720	0	0	44	98%	100%	:f	35%	0	441	256	42204	17127	11305	1311
76%	227	0	0	6508	16105	5477	35376	42124	0	0	44	97%	47%	:	25%	4	554	335	5998	23542	13475	3334
75%	469	0	0	6498	7294	10796	31704	24	0	0	44	96%	0%	-	18%	0	383	201	6762	22559	5964	4391
70%	191	0	0	5866	5892	2339	18748	8	0	0	44	98%	0%	-	12%	0	439	88	10071	17085	3473	1466
80%	329	0	0	5835	7431	1264	22688	35268	0	0	44	97%	41%	Ff	21%	0	379	103	9808	22145	5685	524
81%	376	0	0	6494	10551	1731	27178	94209	0	0	44	99%	100%	:f	30%	0	586	171	40824	19363	8734	743
73%	351	0	0	5863	15789	2611	14771	28223	0	0	44	99%	33%	:	17%	2	440	258	17085	15384	11795	1710

The top frame of Figure 12 shows the load on storage controller A. The controller has a large amount of FC IOPS because it supports the Oracle and SQL Server Databases as well as the Exchange Server log files. The amount of iSCSI and NFS IOPS is smaller because the controller supports 25 VDI desktops using NFS and the SharePoint log files using iSCSI.

The bottom frame of Figure 12 shows the load on storage controller B. The amount of FC IOPS is smaller because this controller supports the Oracle and SQL Server log files and the Exchange Server Database files using FC. In addition, the amount of iSCSI and NFS IOPS is also smaller because the controller supports 25 VDI desktops using NFS and the SharePoint database files using iSCSI.

5 CONCLUSION

With the introduction of Unified Connect, NetApp significantly enhances the value proposition of the converged network by allowing our unified storage products, in conjunction with FCoE–enabled 10GbE switches, to provide support for the concurrent transmission of FCoE, iSCSI, NFS, and CIFS data over a common shared 10GbE target adapter. Consistent with the NetApp Unified Storage Architecture, Unified Connect offers the following benefits:

- True end-to-end network convergence including IP and FCoE storage traffic
- Reduced overall complexity by reducing the number of management points at the target
- Improved efficiency by increasing the overall bandwidth utilization, which allows a more effective use of your data center infrastructure
- Reduced cabling requirements and significantly improved port utilization
- Agile and high-performing network for increased flexibility and lower CAPEX

Enjoy the benefits of simplified data management, investment protection, and reduced total cost of ownership by deploying our leading-edge Ethernet storage solutions.

6 APPENDIXES

6.1 HARDWARE RESOURCES

The following equipment was used to validate the configuration.

Table 1) Hardware configuration.

Solution Component	Notes
Primary Storage	
One NetApp FAS3160 with active-active storage controllers	Data ONTAP 8.0.1RC1
2 x DS4243 disk shelves (24 disks each)	Each disk 450GB/15K/SAS
One 10GbE network adapter per controller supporting Unified Connect	
Networking	
Two Cisco Nexus 5010 switches	
Four VM Hosts (ESX, Hyper-V)	

Solution Component	Notes
HP ProLiant DL360 G6	Two hosts were configured for ESX and two for Hyper-V
32GB RAM per host	
Two quad-core Xeon 2.133Ghz processors	
One dual-port QLogic 8152 10Gb CNA	
AIX Machine	
IBM P520	
4GB RAM	
Two-way Power 6 4.2Ghz processor	
One dual-port 8Gb HBA	

6.2 SOFTWARE RESOURCES

The following software components were used to validate the configuration.

Table 2) Software configuration.

Solution Component	Minimum Revision
Primary Storage	
Data ONTAP	8.0.1RC1
FCP, NFS, iSCSI, FlexClone®, SnapDrive for Windows licenses	N/A
NetApp Management Software	
NetApp SnapDrive	6.3
VMware vSphere Infrastructure	
ESX hosts	VMware ESX 4.1
vCenter Server	4.1
View Connection Broker	
View Client	
Hyper-V Infrastructure	
Hyper-V hosts	Windows Server 2008 R2

Solution Component	Minimum Revision
AIX Infrastructure	
AIX host	AIX 6.1 TL5 SP2
Microsoft Applications	
Microsoft Exchange Server	2010 Enterprise Edition
Microsoft Office SharePoint Server	2010 Enterprise Edition
Microsoft SQL Server	2008 Enterprise Edition

7 ACKNOWLEDGMENTS

We thank the following people for their contributions:

- Toby Creek (NetApp)
- Bhavik Desai (NetApp)
- Chris Gebhardt (NetApp)
- Saad Jafri (NetApp)
- Abhinav Joshi (NetApp)
- Patrick Strick (NetApp)

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

www.netapp.com

Copyright 2010 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexClone, FlexVol, and SnapDrive are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Cisco is a registered trademark of Cisco Systems, Inc. Xeon is a registered trademark of Intel Corporation. Microsoft, Exchange, Internet Explorer, SQL Server, SharePoint, Windows, Windows XP, and Windows Media are registered trademarks and Hyper-V is a trademark of Microsoft Corporation. Oracle is a registered trademark and Oracle11*g* is a trademark of Oracle Corporation. VMware is a registered trademark and vSphere and vCenter are trademarks of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3875-1010