



Dell EqualLogic Best Practices Series

Sizing and Best Practices for Deploying Microsoft Exchange Server 2010 on VMware vSphere and Dell EqualLogic Storage

A Dell Technical Whitepaper

Ananda Sankaran

Storage Infrastructure and Solutions

Dell Product Group
July 2010

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2010 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, and the *DELL* badge, *PowerConnect*[™], *Dell*[™] *EqualLogic*[™], *PowerEdge*[™] and *PowerVault*[™] are trademarks of Dell Inc. *Broadcom*[™] is a registered trademark of Broadcom Corporation. *Intel*[™] is a registered trademark of Intel Corporation in the U.S. and other countries. *Microsoft*[®][™], *Windows*[™], *Windows Server*[™], and *Active Directory*[™] are either trademarks or registered trademarks of Microsoft[®] Corporation in the United States and/or other countries.

Table of Contents

1	Introduction.....	1
1.1	Audience.....	1
1.2	Using an EqualLogic SAN in a Virtualized Exchange Environment	1
1.3	The Rest of this Paper.....	2
2	Mailbox Server I/O Profiling	3
2.1	Test System Configuration.....	4
2.1.1	VMware ESX Host Configuration	5
2.2	Baseline I/O Profile Study using Windows 2008 R2 iSCSI Initiator	11
2.3	Array Type Study using Windows 2008 R2 iSCSI Initiator	14
2.4	RAID Type Study using Windows 2008 R2 iSCSI Initiator	16
2.5	Database Volume Layout Study using Windows 2008 R2 iSCSI Initiator	18
2.6	SAN Scalability Study using Windows 2008 R2 iSCSI Initiator	20
3	Sizing EqualLogic Storage for Exchange Server 2010 Mailbox Role in a Virtualized Environment..	23
3.1	Exchange Random IOPS and EqualLogic Storage.....	23
3.1.1	Disk Drive Rotational Speed.....	23
3.1.2	Number of Disk Drives in the Array	23
3.1.3	RAID Policy	23
3.2	EqualLogic Storage – Usable Capacity	24
3.3	Sizing EqualLogic Storage for Microsoft Exchange.....	25
3.3.1	Microsoft Guidelines for Storage Capacity Sizing.....	25
3.3.2	Microsoft Guidelines for Storage Performance Sizing.....	26
3.3.3	Database Capacity Sizing Process.....	27
3.3.4	Log Capacity Sizing Process	28
3.3.5	Database I/O Performance Sizing Process.....	30
3.3.6	Log I/O Performance Sizing Process.....	30
3.4	Choosing EqualLogic Array.....	31
3.4.1	Selection Process	31
3.4.2	Sizing Example for an Exchange 2010 Environment Supporting 5000 users	32
4	Exchange 2010 Database Availability Group (DAG) Studies.....	34
4.1	Workload Simulation and Server Performance Criteria	34

4.2	Test Configuration	34
4.3	Active/Passive Configuration	37
4.3.1	Active/Passive Configuration Test Results	37
4.4	Active/Active Configuration	39
4.4.1	Active/Active Configuration Test Results	40
5	Best Practice Recommendations	43
5.1	Storage	43
5.2	SAN Design	44
5.3	Configuring Host Server Resources	45
5.4	ESX Host Configuration	46
5.5	EqualLogic Multipathing Extension Module for VMware vSphere 4.1	47
5.6	Virtual Machine and Guest OS Configuration	48
5.7	Exchange Server Configuration	49
Appendix A	Test System Components	50

1 Introduction

Many organizations will be looking to move to Microsoft® Exchange 2010 to achieve lower TCO while taking advantage of new capabilities that enhance performance, reliability and scalability of the platform. Key innovations in the core Exchange 2010 platform focus on providing enhanced mailbox resiliency. To achieve optimal use of new Exchange 2010 mailbox resiliency features such as Database Availability Groups and online mailbox moves, it is critical that your database storage infrastructure be properly designed and sized to meet storage and I/O performance requirements.

IT organizations are also continuing to improve the efficiency and availability of enterprise applications through the use of datacenter virtualization architectures. With the release of VMware® vSphere™ 4.1, VMware has included new functional enhancements that improve integration and performance of Storage Area Networks (SANs).

The goal of this paper is to present the results of a series of Exchange 2010 storage I/O performance tests conducted by Dell Labs, and provide storage sizing guidance and best practices based on those results for designing and running Microsoft® Exchange 2010 hosted by VMware vSphere 4.1 running on Dell PowerEdge blade servers and Dell™ EqualLogic™ SAN. This paper also provides generalized storage sizing formulas for planning storage requirements, and an example reference configuration for supporting a 5000 user Exchange environment.

1.1 Audience

This paper is intended for Solution Architects, Storage Network Engineers, System Administrators and IT Managers who need to understand how to design, properly size and deploy Microsoft® Exchange Server 2010 running on the VMware vSphere virtualization platform and Dell EqualLogic storage. It is expected that the reader has working knowledge of Microsoft® Exchange deployment and operation, VMware system administration and EqualLogic based iSCSI SAN operation.

1.2 Using an EqualLogic SAN in a Virtualized Exchange Environment

Using an EqualLogic SAN as your storage infrastructure in virtualized Exchange environments provides multiple benefits:

- Dynamic Infrastructure** EqualLogic PS Series SAN coupled with Dell PowerEdge™ servers running VMware ESX Server enable quick provisioning, high performance, scalability, and the flexibility to dynamically deploy server, storage and application stacks.
- Seamless Scalability** SAN storage can easily be expanded to accommodate for increased capacity or performance needs. Similarly, hardware resources and operating systems for Exchange servers can be quickly deployed in the virtualized server infrastructure as needed.

Data Protection	EqualLogic SAN features such as volume cloning and Exchange aware snapshots with Auto Snapshot Manager/Microsoft® Edition (ASM/ME) can provide comprehensive protection and granular recovery of Exchange data at no additional cost.
Efficient Management	Instead of managing disconnected pools of Direct Attached Storage (DAS), an EqualLogic SAN can be managed as one single entity with pools of data allocated for specific purposes. With virtualized storage and server infrastructure, there are fewer physical assets to manage, hence overall cost of managing and supporting the physical infrastructure is reduced over time.
Path Redundancy	EqualLogic Multi-Path I/O software for Microsoft® Windows Server 2008 and the Advanced Multipathing Extension Module (MEM) for VMware vSphere 4.1 allow you to take full advantage of multiple network adapters to provide network path redundancy and I/O load balancing.
High Availability	Using built-in HA and DR features of Exchange 2010, customers can improve availability of their Exchange environment in case of a local or a site failure. By encapsulating exchange roles in virtual machines, dissimilar hardware can be used at DR sites to support failover from production sites. Redundant, hot serviceable components in all of the infrastructure solution components improve component level availability, thus maximizing application uptime.

1.3 The Rest of this Paper

The rest of this paper contains the following sections:

- **Section 2, Mailbox Server I/O Profiling** on page 3
- **Section 3, Sizing EqualLogic Storage for Exchange Server 2010 Mailbox Role in a Virtualized Environment** on page 23
- **Section 4, Exchange 2010 Database Availability Group (DAG) Studies**, on page 34
- **Section 5, Best Practice Recommendations**, on page 43
- **Appendix A ,Test System Components**, on page 50

2 Mailbox Server I/O Profiling

To support development of this paper we deployed and tested a full-scale virtualized Microsoft® Exchange server configuration running within a VMware vSphere v4.1 virtual machine environment at Dell Labs.

In the first phase of our lab tests we conducted a series of I/O profiling studies to establish baseline Exchange Server 2010 mailbox server I/O load profiles. We used Jetstress¹ to simulate Exchange I/O workloads.

Note: Jetstress is typically used to validate performance and stability of I/O storage subsystems before putting Exchange systems into production. Jetstress uses Exchange database (ESE – Extensible Storage Engine) binaries to simulate Exchange I/O. For more information on Jetstress, see:

➤ *Microsoft® Exchange Server Jetstress 2010:*

<http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=13267027-8120-48ed-931b-29eb0aa52aa6>

In all test cases, Jetstress was run with the following load attributes:

- Background database maintenance was turned on while simulating the mailbox users I/O
- Test duration was 2 hours for each simulation run
- Database checksum was executed after completion of each test
- The maximum cache memory used by Jetstress for each database in test was 256MB

Jetstress response time criteria or latency requirements for I/O:

- Database read latency threshold: 20ms
- Log write latency threshold: 10ms

¹ Jetstress 2010 Beta, version 14.00.0639.013 with Exchange 2010 RTM binaries

2.1 Test System Configuration

Configuration details for the major system components in our test system configuration are provided in Appendix A on page 50. The iSCSI SAN network and test system components² are shown in Figure 1.

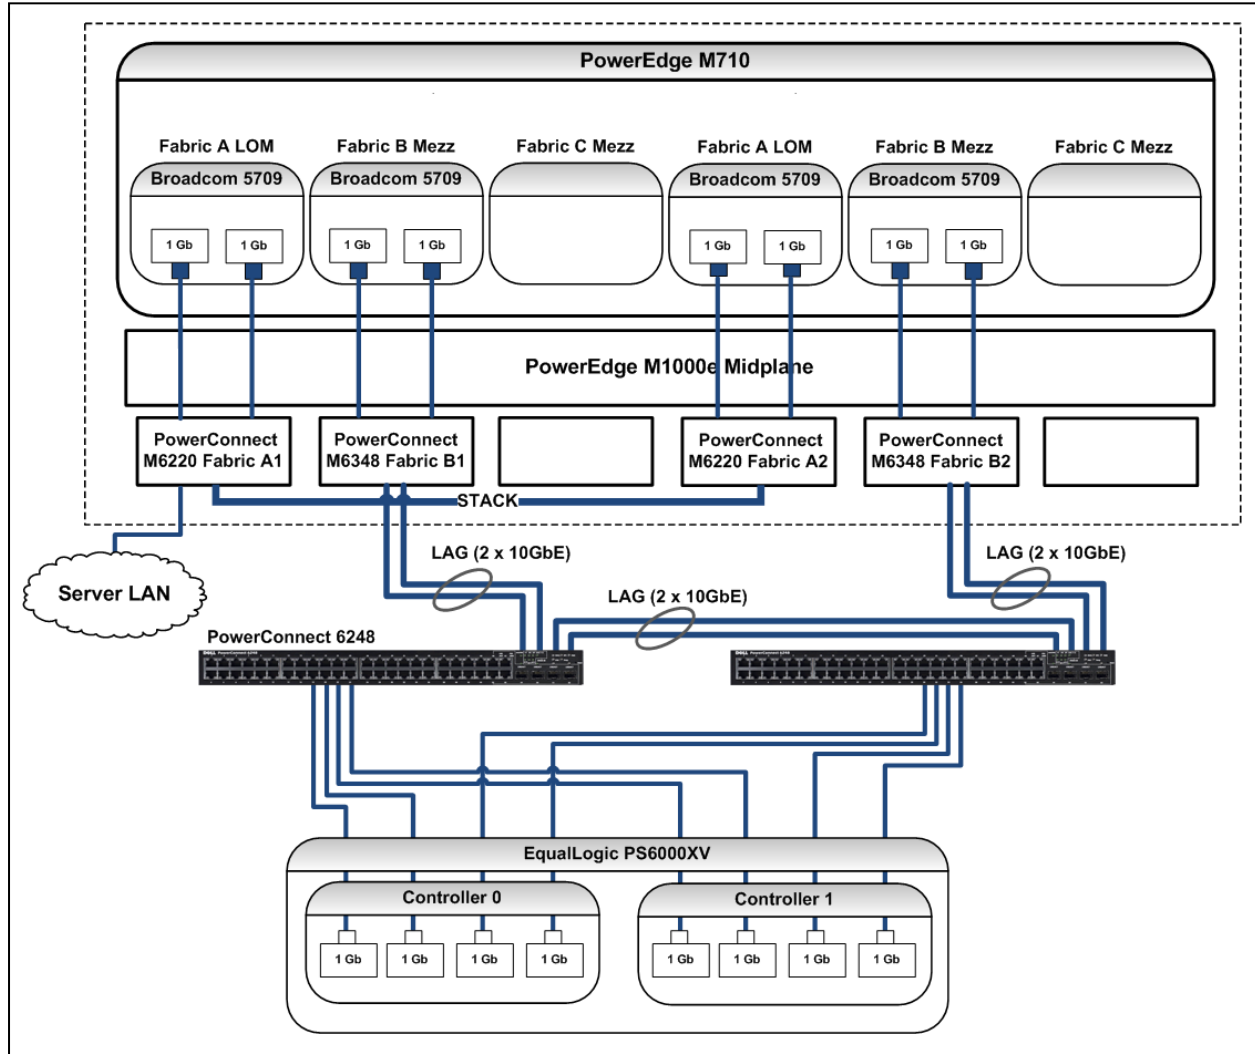


Figure 1 Baseline I/O Profiling Test System Configuration

² Instead of PowerConnect M6348, PowerConnect M6220 switches can also be used for Fabric B.

2.1.1 VMware ESX Host Configuration

VMware ESX Server v4.1 was installed on the Dell PowerEdge M710 server shown in Figure 1. Microsoft® Windows Server 2008 R2 Enterprise Edition was installed in a guest virtual machine running on the ESX host.

ESX host configuration overview:

- The mailbox VM was configured to use four virtual CPUs and 48GB of reserved memory at the ESX hypervisor layer.
- The blade server local storage was configured using the two installed disk drives configured as a RAID 1 set. A portion of this local storage was used for installation of the VMware ESX server software. The remaining local storage was formatted as a VMFS partition.
- The guest VM operating system images were created on the local VMFS data store. The Exchange Server databases and logs were stored in volumes hosted by the externally connected EqualLogic PS6000XV array shown in Figure 1.
- Microsoft® software iSCSI initiator installed within the Exchange VM is used to directly access the database and log volumes on the EqualLogic SAN. Direct access of Exchange data and log volumes using the iSCSI initiator within the VM enables using EqualLogic features such as Exchange consistent snapshots and clones and Exchange protection tools such as Auto-Snapshot Manager (ASM)/Microsoft® Edition. Using these features is not possible when Exchange data and log volumes are exposed as VMware virtual disks (using the software or hardware initiator at the VMware ESX layer).

Note: For more details describing Exchange 2010 support requirements for Hardware Virtualized environments, see the following Microsoft® TechNet article:

- *Exchange 2010 System Requirements:* <http://technet.microsoft.com/en-us/library/aa996719.aspx>

Physical Networking Details

In this section we describe the physical configuration of the network components shown in Figure 1 in more detail.

- Four Fabric A NIC ports on the M710 server are connected to the server LAN via a pair of stacked PowerConnect M6220 switch modules installed in the M1000e blade chassis (I/O modules A1 and A2).
- Four Fabric B NIC ports on the M710 server are connected to the iSCSI SAN via a pair of PowerConnect M6348 switch modules installed in the M1000e blade chassis (I/O modules B1 and B2). (Note: In this test configuration PowerConnect M6220 or M6348 switches could have been used in both Fabric A and Fabric B)
- Each of the M6348 switches is connected to the top-of-rack PowerConnect 6248 switches via Link Aggregation Groups (LAG). The LAGs combined two 10GbE uplinks to create 20Gbps total bandwidth using the uplink modules on respective switches.
- The two top-of-rack 6248 switches are also connected via a 20Gbps LAG group using a second uplink module installed in each switch.

Table 1 and Table 2 show the functional assignment of blade server switch fabric modules and server NIC port assignments.

M1000e Chassis Switch Module	Switch Module	Purpose
I/O Module A1	PowerConnect M6220	Server LAN
I/O Module A2	PowerConnect M6220	Server LAN
I/O Module B1	PowerConnect M6348	iSCSI Storage
I/O Module B2	PowerConnect M6348	iSCSI Storage

Table 1 Blade Server Fabric Switch Assignments

Blade Server	I/O Module	Number of NIC Ports (via a Mezzanine card)	Purpose
PowerEdge M710	I/O Module A1	2	Server LAN
	I/O Module A2	2	
	(Total ports on Fabric A = 4)		
	I/O Module B1	2	iSCSI Storage
	I/O Module B2	2	
	(Total ports on Fabric B = 4)		

Table 2 Blade Server NIC Port Assignment Details

ESX Virtual Network Details

On the ESX host, the virtual switches were configured as follows:

- One virtual switch was configured for iSCSI SAN access, with the four M710 NIC ports on Fabric B as uplinks.
- One virtual switch was configured for server LAN access with the four M710 NIC ports on Fabric A as uplinks.

The vSwitch configuration on ESX host and blade server NIC port assignments to those vSwitches is listed in tables below. It is recommended to enable jumbo frames on the host vSwitch. The Broadcom 5709 NICs installed on blade servers support jumbo frames. Note: In ESX 4.1, Jumbo Frames are not supported when using hardware iSCSI initiators at the ESX layer.

vSwitch	vmnics assigned	Purpose
vSwitch0	vmnic0, vmnic1, vmnic2, vmnic3	Server LAN
vSwitch1	vmnic6, vmnic7, vmnic8, vmnic9	iSCSI Storage
<p>Note: vmnics 4, 5, 11 and 12 connected to fabric C were not used in test. More information on vmnic enumeration within vSphere hosts on the Dell PowerEdge M1000e can be found here: http://content.dell.com/us/en/enterprise/d/business~solutions~engineering-docs~en/Documents~NetworkingGuide_vSphere4_Blades.pdf.aspx</p>		

Table 3 vSwitch and vmnic assignments

As a best practice, different kinds of network traffic should be physically or logically segregated. In our environment, we segregate network traffic by using VLANs for different classes of traffic (tagged packets) at the virtual switch layer³. VLAN and port group assignments (on vSwitch0, the Server LAN) for the vSphere host were assigned as shown in Table 4.

VLAN ID	Port Group Label	Port Group Type	Purpose
101	Service Console	Service Console	Service console management traffic
102	VMkernel	VMkernel	Kernel level traffic (e.g. vMotion)
103	VM Network – Public	Virtual Machine	Guest VM traffic (Exchange client/server LAN)
104	VM Network – Private	Virtual Machine	Guest VM traffic (ie. Exchange server private replication)

Table 4 vSwitch0 VLAN and Port Group Assignments

³ This does not apply to the iSCSI storage LAN (vSwitch1 in our configuration).

The vSwitch0 configuration on the ESX host is shown in Figure 2.

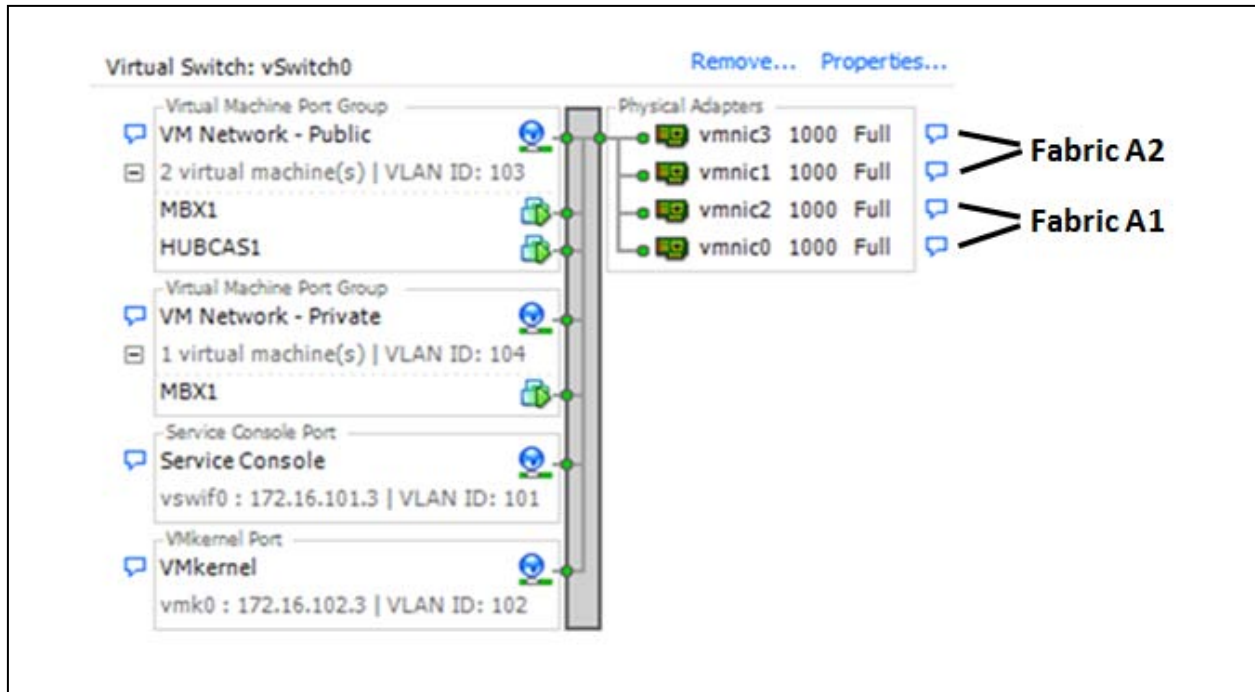


Figure 2 vSwitch0 Configuration

The vSwitch1 configuration on the ESX host is shown in Figure 3⁴. The iSCSI initiator in the Exchange mailbox VM was used to access the database and log volumes on the EqualLogic SAN in this configuration.

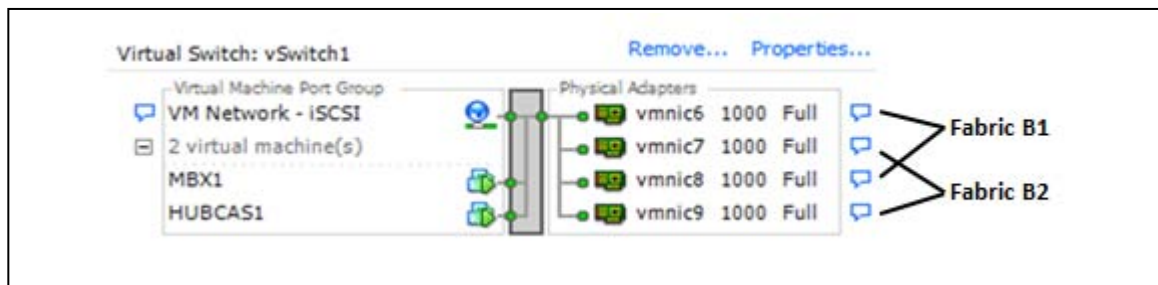


Figure 3 vSwitch1 Configuration

⁴ Only virtual machine MBX1 was used in the Jetstress studies. HUBCAS1 VM was shut down.

Note: The load balancing policy was configured at the vSwitch for load balancing across the multiple NIC uplinks. The default policy was used, which is 'Route based on the originating virtual switch port ID'. For further information, see the following VMware Information Guide article:

- *VMware Virtual Networking Concepts*,
http://www.vmware.com/files/pdf/virtual_networking_concepts.pdf

Virtual Machine NIC Assignments and Networking

The virtual network adapter assignments for the mailbox server VM (using the Windows 2008 R2 iSCSI initiator within the VM) are as follows:

- The native OS iSCSI initiator within the guest OS VM was used to connect the data volumes on EqualLogic iSCSI storage.
- Two virtual NICs were configured within the mailbox server VM connecting to the virtual switch (vSwitch0) configured for server LAN access on respective VLANs.
- Four virtual NICs were configured within the mailbox server VM connecting to the virtual switch (vSwitch1) configured for storage.
- The EqualLogic Host Integration Toolkit⁵ (HIT) was installed in the VM to provide the EqualLogic DSM (Device Specific Module) and enable MPIO (Multi-path I/O) to storage.

Table 5 shows the virtual machine NIC assignments.

Virtual Network Adapter	Adapter Type	vSwitch	Port Group	VLAN
# 1	E1000	vSwitch0	VM Network - Public	103
# 2			VM Network - Private	104
# 3	VMXNET 3	vSwitch1	VM Network – iSCSI	N/A
# 4				
# 5				
# 6				

Table 5 Exchange mailbox VM NIC Assignments

⁵ The EqualLogic Host Integration Toolkit is available at the following download location (support login ID required): <https://www.equallogic.com/support/download.aspx?id=3199>

The Exchange mailbox virtual machine virtual network adapter and VLAN connectivity described in Table 5 is shown in Figure 4.

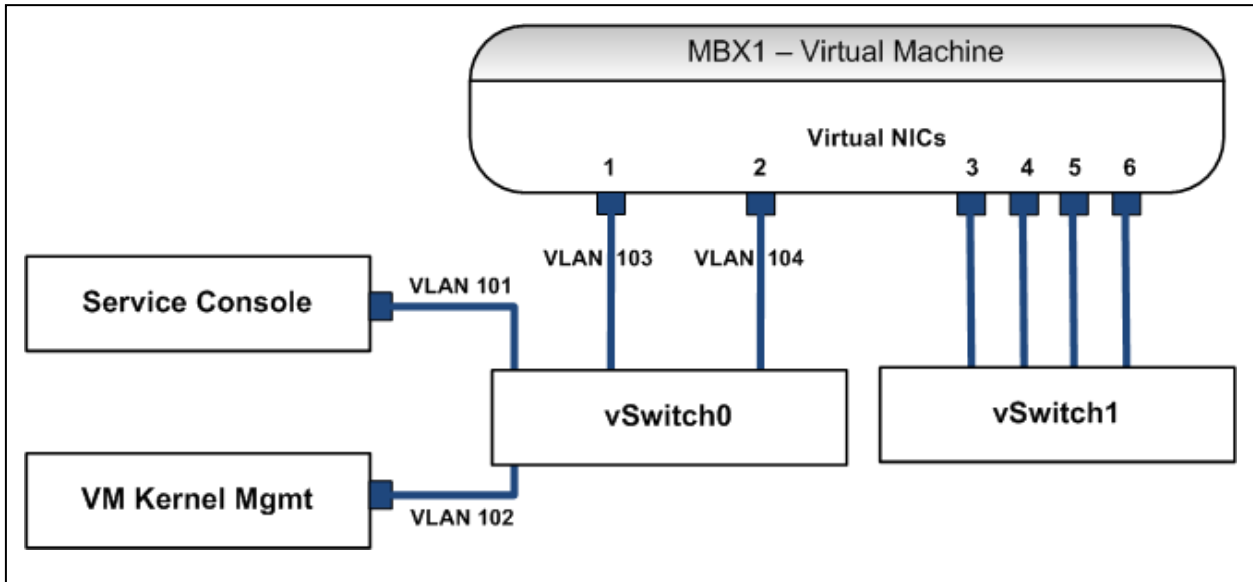


Figure 4 Virtual Network Configuration for the Exchange mailbox VM

Figure 5 shows the Exchange mailbox VM hardware device configuration, including the network adapters for the configuration shown in Figure 4.

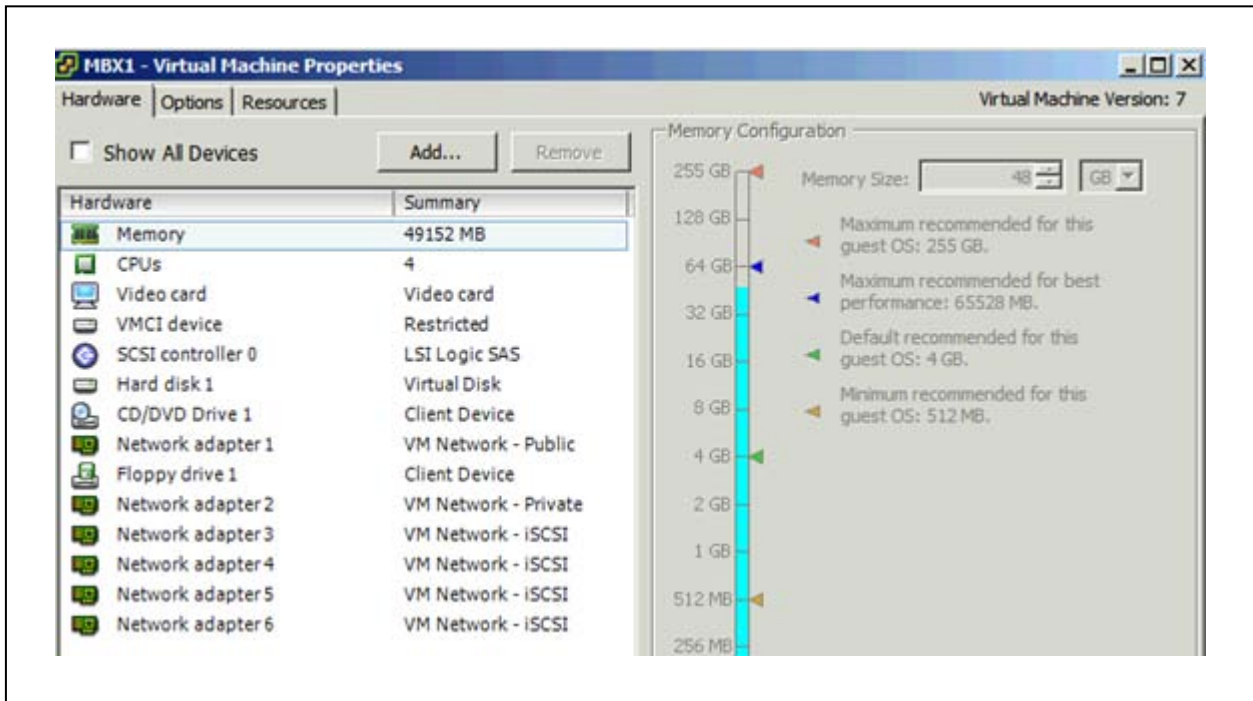


Figure 5 MBX 1 Virtual Machine Properties

2.2 Baseline I/O Profile Study using Windows 2008 R2 iSCSI Initiator

The goal of this test was to establish baseline I/O performance characteristics of the test system configuration in support of Microsoft® Exchange Server 2010 end-user mail operations (via Microsoft® Outlook or other mail clients).

In the baseline I/O profile study, we simulated Exchange Server 2010 user mailboxes with different IOPS requirements per user profile. A single EqualLogic PS6000XV array configured using RAID 50 was used for hosting the database and log volumes. The test system topology is shown in Figure 1 on page 4. The EqualLogic Array configuration used corresponds to that shown for "Volume Layout Test #1" in Table 9 on page 19.

Exchange Server 2010 mail operations generate a certain amount of IOPS (I/O Operations per Second) per operation between server and storage. Microsoft® has provided a TechNet article⁶ which includes tables for estimating Exchange 2010 Mailbox IOPS based on different user profiles. The user profile is defined by how much database cache is used per user, and how many messages each user sends and receives per day.

For example, a user mailbox with 150 messages sent or received per day with 9MB Database cache allocated per user mailbox correlates to 0.18 IOPS if the database is configured in the stand-alone mode, and correlates to 0.15 IOPS if the database is configured in copy mode (Exchange Database Availability Group, or "DAG"). These values correspond to the IOPS generated on the database only, 2GB mailbox size and for an average message size of 75KB. They do not include the IOPS required for other Exchange mailbox server activities such as log writes, log replication and database maintenance. Mailbox databases can have replicated database copies configured if the mailbox servers hosting them are configured as part of a DAG. The servers in a DAG can host copies of mailbox databases from other servers in the DAG. Each server can host multiple databases but only one copy instance per database.

Test parameters for this configuration are shown in Table 6.

⁶ Understanding Database and Log Performance Factors: <http://technet.microsoft.com/en-us/library/ee832791.aspx>

Microsoft Exchange:	
Simulated # of Users	5000
Simulated mailbox size	512MB
Mailbox databases	4 x standalone, 1250 users per database
Database volumes	1024GB each, one per mailbox database (4 total), GPT Disks, 64KB NTFS allocation unit size
Log volumes	150GB each, one per mailbox database), GPT Disks, 64KB NTFS allocation unit size
Storage Array	
PS6000XV: 15K 450Gb SAS disks, 14 disks as RAID 50, 2 x 450GB as hot spares	
A single (default) was pool used	

Table 6 Test Configuration Parameters for I/O Baseline Profile

Test results for two user profiles that met the latency requirements of Jetstress are illustrated in Figure 6. The two user loads that were simulated were for 5000 users running at 0.2 IOPS per user and 0.35 IOPS per user respectively. A 0.2 IOPS per user load corresponds to a user profile of 150 messages sent /received per day with 9MB cache allocated per mailbox if the database is configured in a stand-alone mode, or 200 messages sent / received per day with 12 MB cache allocated per mailbox if database is in copy mode. A 0.35 IOPS per user load corresponds a user profile of almost 300 messages sent / received per day with 18MB cache allocated per mailbox in standalone mode, or 350 messages sent / received per day with 21MB cache allocated per mailbox in the copy mode.

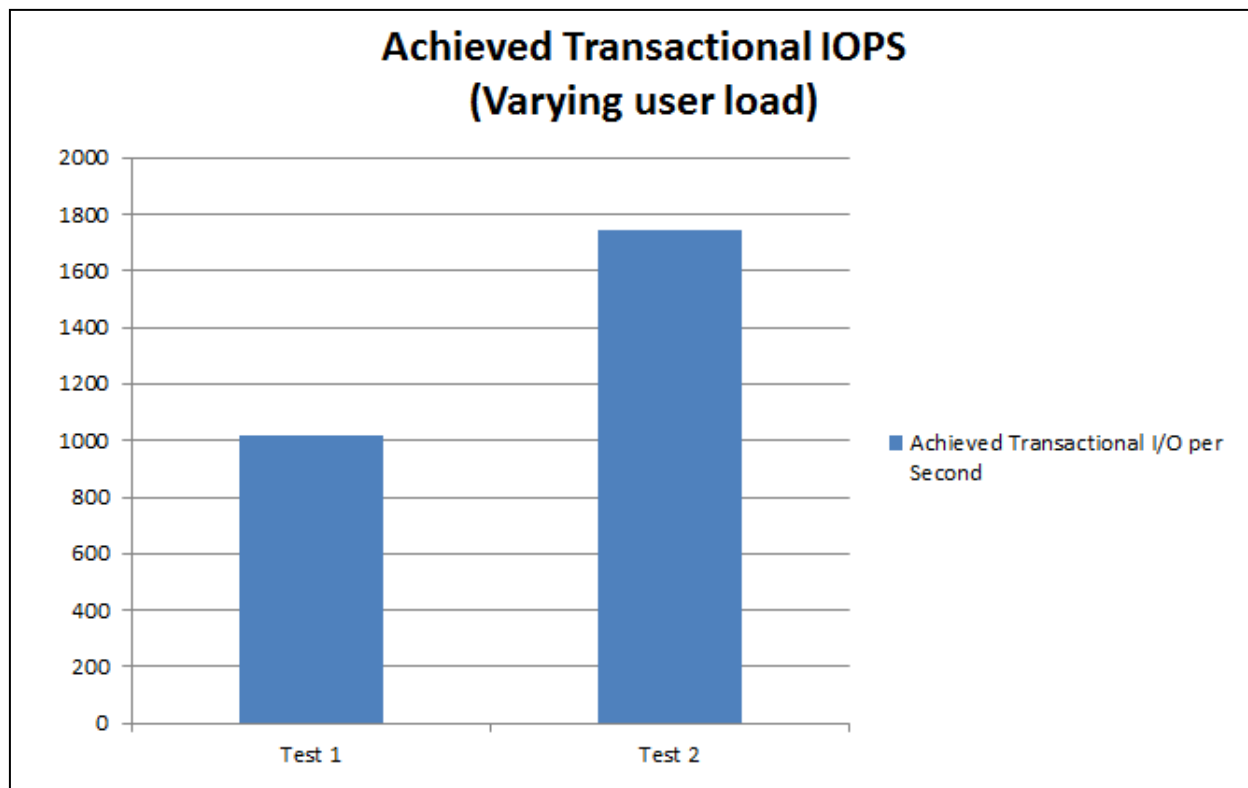


Figure 6 Measured IOPS: Comparison for Two Different Mailbox Profiles (User Loads)

The achieved transactional IOPS includes both read and write database IOPS during the test. Additional IOPS were incurred on storage for log write operations and database background maintenance read activity. These additional IOPS are not included⁷ in the results shown in Figure 6 as these IOPS support achieving the effective transactional I/O. The achieved transactional IOPS is generally used for communicating user profile IOPS. As seen in Figure 6, the average IOPS increases with applied user load.

Database I/O latencies are illustrated in Figure 7 below for each test.

⁷ Database IOPS is the standard metric for sizing mailbox users based on the IOPS per mailbox profile.

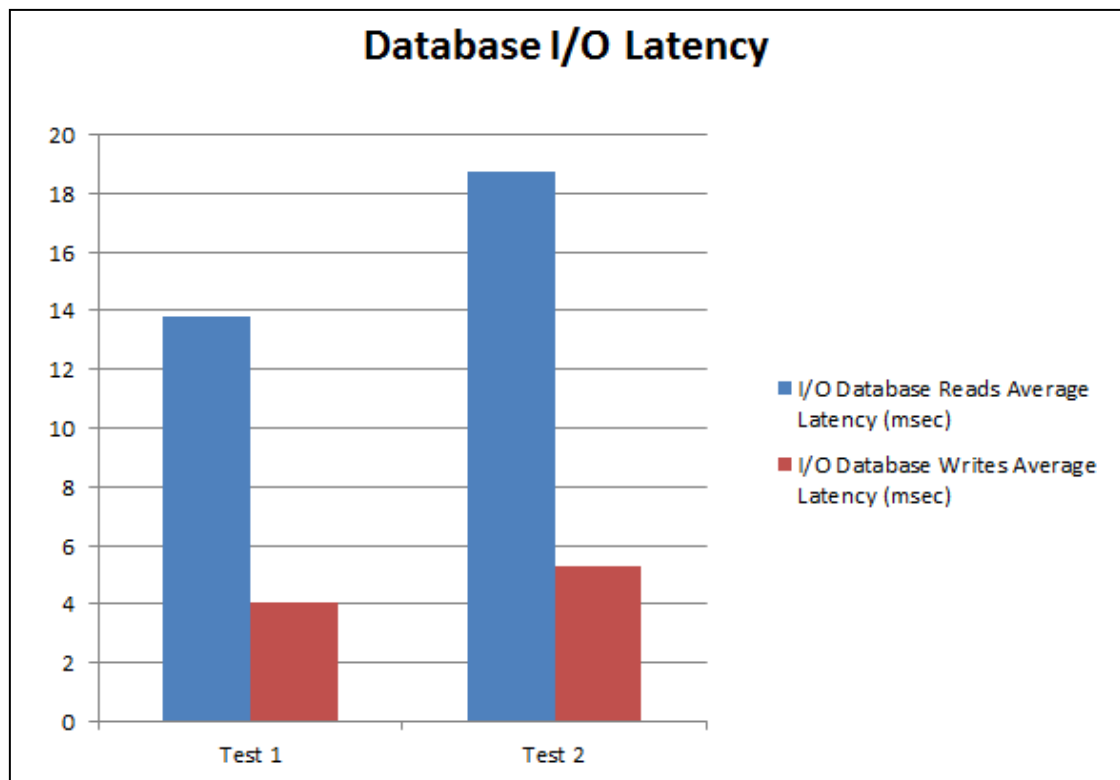


Figure 7 Database I/O Latency for Two Mailbox Profiles or User Loads

As the generated load increases, the total measured IOPS from storage increases. The corresponding latency also increases with IOPS, but within acceptable limits. The results of this test show that the PS6000XV array configured with RAID 50 can provide up to 1750 IOPS within acceptable latency limits for Exchange 2010 hosted on VMware vSphere. This translates to 5000 users with an IOPS profile of 0.35 IOPS per mailbox, or 10000 users with an IOPS profile of 0.17 IOPS per mailbox. The IOPS requirement per user mailbox will depend on the number of messages sent/received per day, database configuration (stand-alone or DAG copy) and the amount of mailbox memory cache allocated per user⁸.

2.3 Array Type Study using Windows 2008 R2 iSCSI Initiator

The goal of this test was to measure variation in storage I/O performance for different EqualLogic PS Series array types. Each array type supports a different disk drive type and drive count. The configurations were identical to the configuration in section 2.2 except for the array type, disk drive type and in the case of the PS6500E the number of disk drives.

The Microsoft® iSCSI initiator was used within the mailbox VM for connecting to the iSCSI volumes on the storage array. Test parameters for this configuration are shown in Table 7.

⁸ The IOPS requirements for a range of user profiles is illustrated in the following Microsoft® TechNet article: <http://technet.microsoft.com/en-us/library/ee832791.aspx>

Microsoft Exchange:	
Simulated # of Users	5000
Simulated IOPS per user mailbox	0.35
Mailbox databases	4 x standalone
Database volumes	1024GB each, one per mailbox database (4 total), GPT Disks, 64KB NTFS allocation unit size
Log volumes	150GB each, one per mailbox database (4 total), GPT Disks, 64KB NTFS allocation unit size
Storage Arrays	
Test #1:	<u>PS6000XV:</u> <ul style="list-style-type: none"> • 15K 450GB SAS disks • 14 disks as RAID 50 • 2 disks as hot spares
Test #2:	<u>PS6500E:</u> <ul style="list-style-type: none"> • 7.2K 1024GB SATA II disks • 46 disks as RAID 50 • 2 disks as hot spares
Test #3:	<u>PS6000X:</u> <ul style="list-style-type: none"> • 10K 600GB SAS disks • 14 disks as RAID 50 • 2 disks as hot spares

Table 7 Test Configuration Parameters for Array Type Study

The average database IOPS obtained with the different array types are shown in Figure 8. The disk latency response times stayed within a 17ms - 20ms range for all configurations.

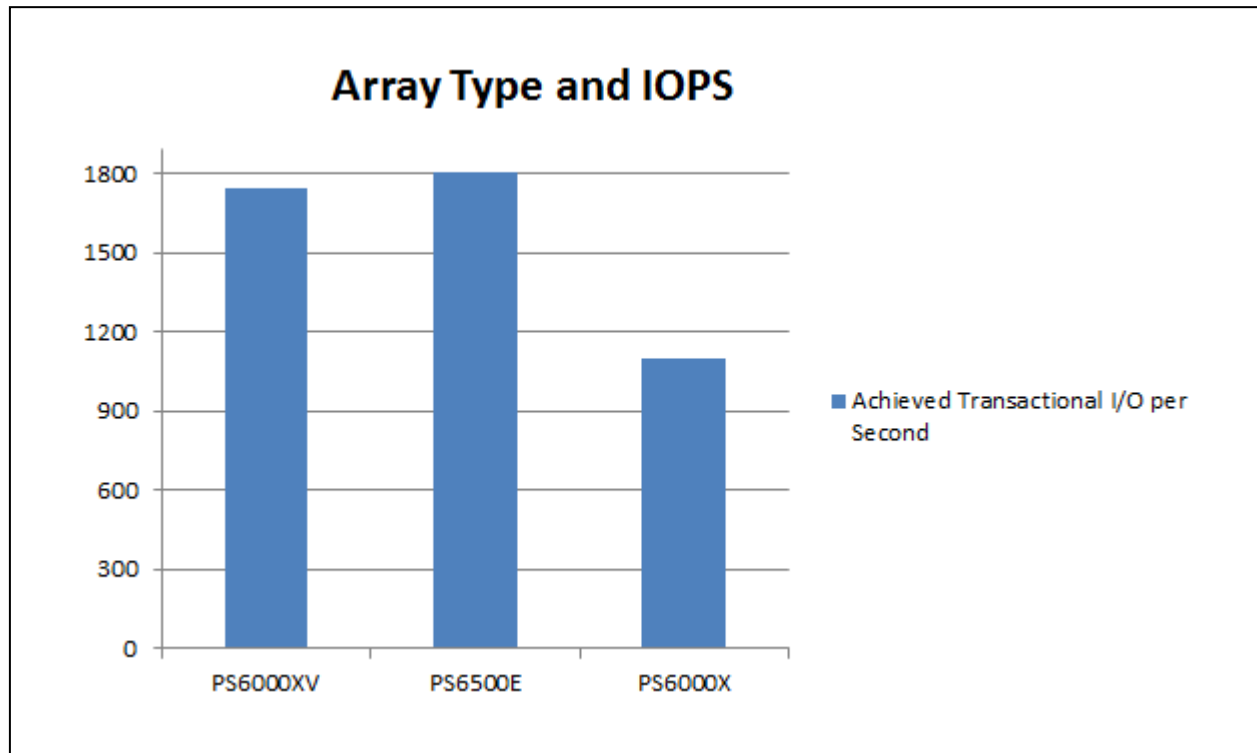


Figure 8 Array Types Study - Average Database IOPS

As shown in Figure 8, the PS6000XV array provides more IOPS than the PS6000X array. This is due to use of the 15K RPM SAS drives, as compared to the 10K RPM SAS drives in the PS6000X. The PS6500E array with SATA II drives produced similar IOPS as the PS6000XV. This was due to the increased drive count in the PS6500E (46 x 7.2K SATA II drives, two hot spares) as compared to the PS6000XV (14 x 15K SAS drives, two hot spares). Random IOPS performance over time depends on the rotational speed of the available drives. In general, higher RPM drives yield more IOPS.

2.4 RAID Type Study using Windows 2008 R2 iSCSI Initiator

The goal of this test was to measure the difference in storage I/O performance when comparing RAID 10 to RAID 50 volumes. In Test #1, a single EqualLogic PS6000XV array configured using RAID 10 hosted both the database and log volumes. In Test #2 a single EqualLogic PS6500E array configured using RAID 10 hosted both the database and log volumes. The configuration used for these tests is identical to the configuration used for the Baseline I/O Profile Study (Section 2.2), with the following exceptions:

- RAID 10 is used instead of RAID 50
- Database volume sizes are 650GB each instead of 1024GB each on the PS6000XV array

The average database IOPS obtained using both RAID 10 and RAID 50 are shown in Figure 9 below. The disk latency response times stayed within a 17ms - 20ms range for both configurations.

- Using the PS6000XV, the RAID 10 configuration provided up to 2166 IOPS and RAID 50 provided up to 1750 IOPS. RAID 10 provided approximately 25% greater IOPS for the same user load. RAID 50 incurs additional disk I/O overhead for parity calculations and additional reads/writes to disk (4 disk I/O operations) to support a single application write operation. In comparison, RAID 10 just mirrors an application write operation to a second disk drive (generating an additional 2 disk I/O operations). Using the PS6500E, RAID 10 supported up to 2414 IOPS and RAID 50 supported up to 1805 IOPS. RAID 10 provided approximately 33% greater IOPS for the same user load.
- Using the same EqualLogic PS6000XV array, we were able to configure 1024GB database volumes in the RAID 50 test, and smaller 650GB database volume sizes for the RAID 10 test. This was due to the fact that, for a given volume size, the data mirroring policy in RAID 10 will consume more physical disk space than the parity data policy implemented in RAID 50.

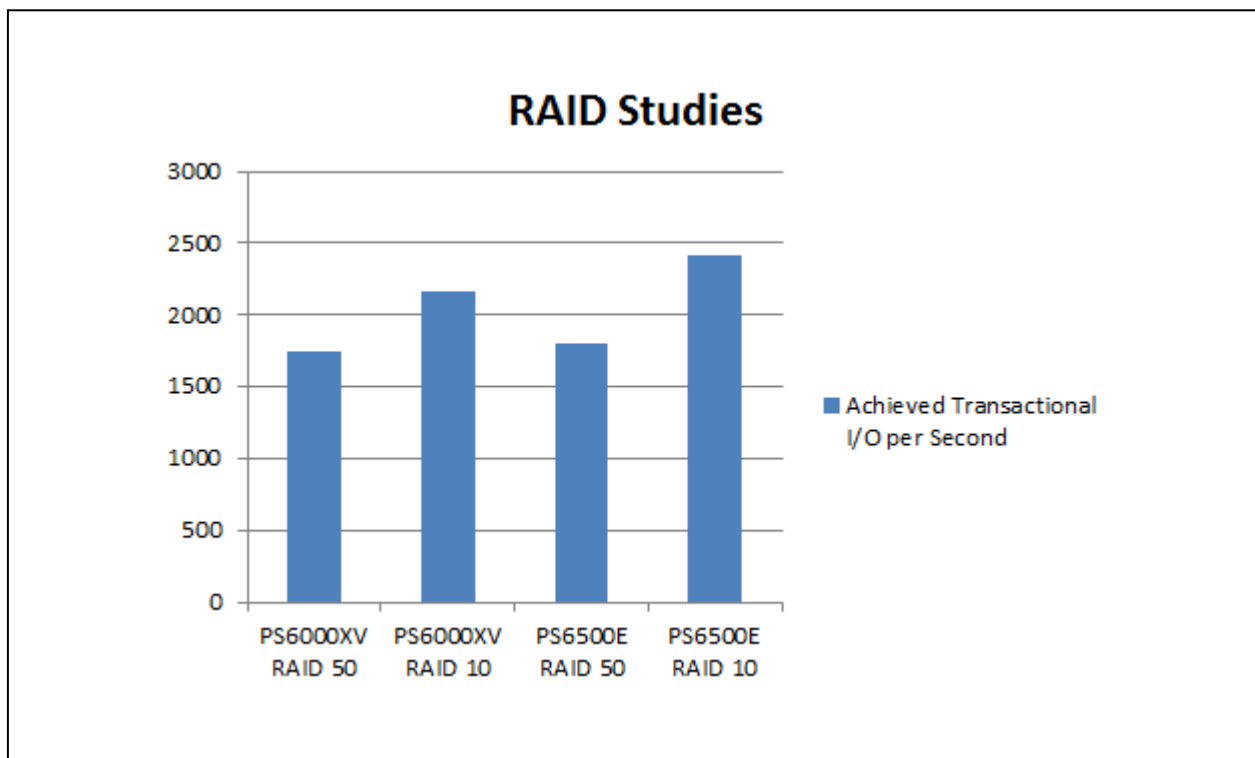


Figure 9 RAID 10 to RAID 50 Transactional IOPS Comparison

2.5 Database Volume Layout Study using Windows 2008 R2 iSCSI Initiator

The goal of this test was to measure variation in storage I/O performance for different database volume and log volume placement strategies. A single EqualLogic PS6000XV array configured with RAID 50 was used for hosting both the database and log volumes. The physical test configuration used was identical to the configuration used in the Baseline I/O Profile test (Section 2.2).

Exchange mailbox test parameters for this configuration are shown in Table 8 below.

Microsoft Exchange:	
Simulated # of Users	5000
Simulated IOPS per user mailbox	0.35
Mailbox databases	4 x standalone
Database volumes	See Table 9
Log volumes	See Table 9
Storage Array	
PS6000XV: 15K 450GB SAS disks, 14 disks as RAID 50, 2 x 450GB as hot spares	
A single (default) pool was used.	

Table 8 Test Configuration Parameters for Database Volume Layout Study

We tested four variations for placement and size of database and log volumes. These test variations are shown in Table 9.

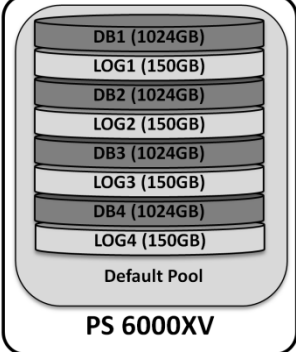
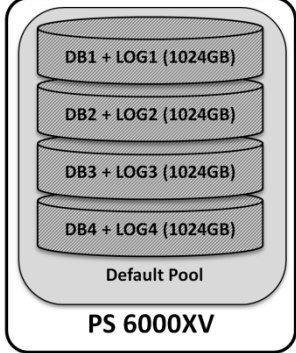
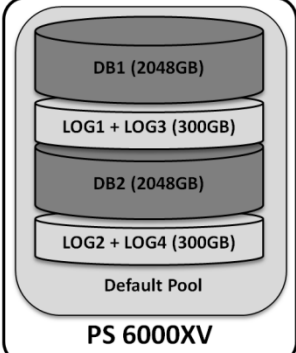
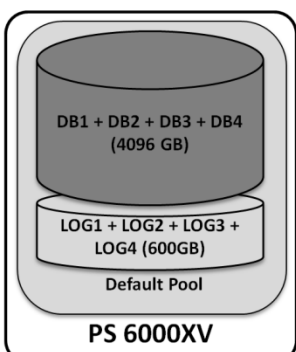
<p>Volume Layout Test #1</p> <ul style="list-style-type: none"> • 4 database volumes / 1250 users per volume (1024GB each) • 4 log volumes (150GB each) • GPT Disks, 64KB NTFS allocation unit size 	
<p>Volume Layout Test #2</p> <ul style="list-style-type: none"> • 4 volumes / 1250 users per volume (1024GB each) • Databases and Logs stored in same volumes • GPT Disks, 64KB NTFS allocation unit size 	
<p>Volume Layout Test #3</p> <ul style="list-style-type: none"> • 2 database volumes / 2500 users per volume (2048GB each) • 2 log volumes (300 GB each) • GPT Disks, 64KB NTFS allocation unit size 	
<p>Volume Layout Test #4</p> <ul style="list-style-type: none"> • 1 database volume/ 5000 users per volume (4096 GB) • 1 log volume (600GB) • GPT Disks, 64KB NTFS allocation unit size 	

Table 9 Volume Layout Study Test Variations

The average database IOPS obtained for each of the volume layout test configurations are shown in Figure 10. The disk latency response times stayed within a 17ms - 20ms range.

All four volume layout configurations achieved IOPS in the same range (all results within 5%). It can be inferred that for a given number of databases, the placement within one or more volumes will not cause noticeable changes in average IOPS.

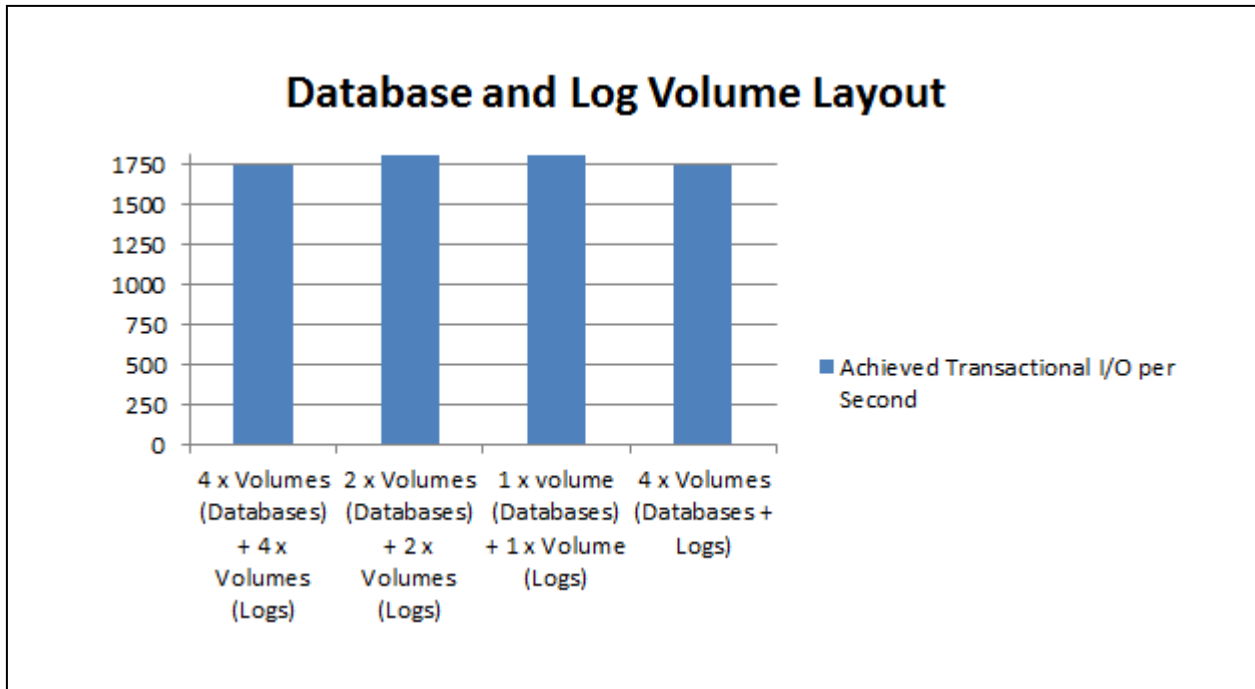


Figure 10 Volume Layout Study - Average Database IOPS

2.6 SAN Scalability Study using Windows 2008 R2 iSCSI Initiator

The goal of this test was to measure how I/O performance scales as you increase the number of EqualLogic PS Series storage arrays (members) with a group. The configurations tested included one, two and three EqualLogic PS6000XV members with the group. User mailbox transactional I/O load was increased linearly during the tests.

Test parameters for this configuration are shown in Table 10.

Microsoft Exchange:	
Simulated # of Users	5000
Simulated IOPS per user mailbox	0.35
Mailbox databases	4 x standalone
Storage Arrays:	
Disk configuration for each PS6000XV: 15K 450GB SAS disks, 14 disks as RAID 50, 2 hot spares	
Test #1:	<ul style="list-style-type: none"> • 1 x PS6000XV • 5,000 users • Database volumes: 4 x 1024GB each • Log volumes: 4 x 150GB each
Test #2:	<ul style="list-style-type: none"> • 2 x PS6000XV (in one EqualLogic storage pool) • 10,000 users • Database volumes: 4 x 2048GB each • Log volumes: 4 x 300GB each
Test #3:	<ul style="list-style-type: none"> • 3 x PS6000XV (in one EqualLogic storage pool) • 15,000 users • Database volumes: 4 x 3072GB each • Log volumes: 4 x 450GB each

Table 10 Test Configuration Parameters for Array Scalability Study

The average database IOPS obtained with the different array quantities is shown in Figure 11. The disk latency response times stayed within a 17ms - 20ms range for all configurations.

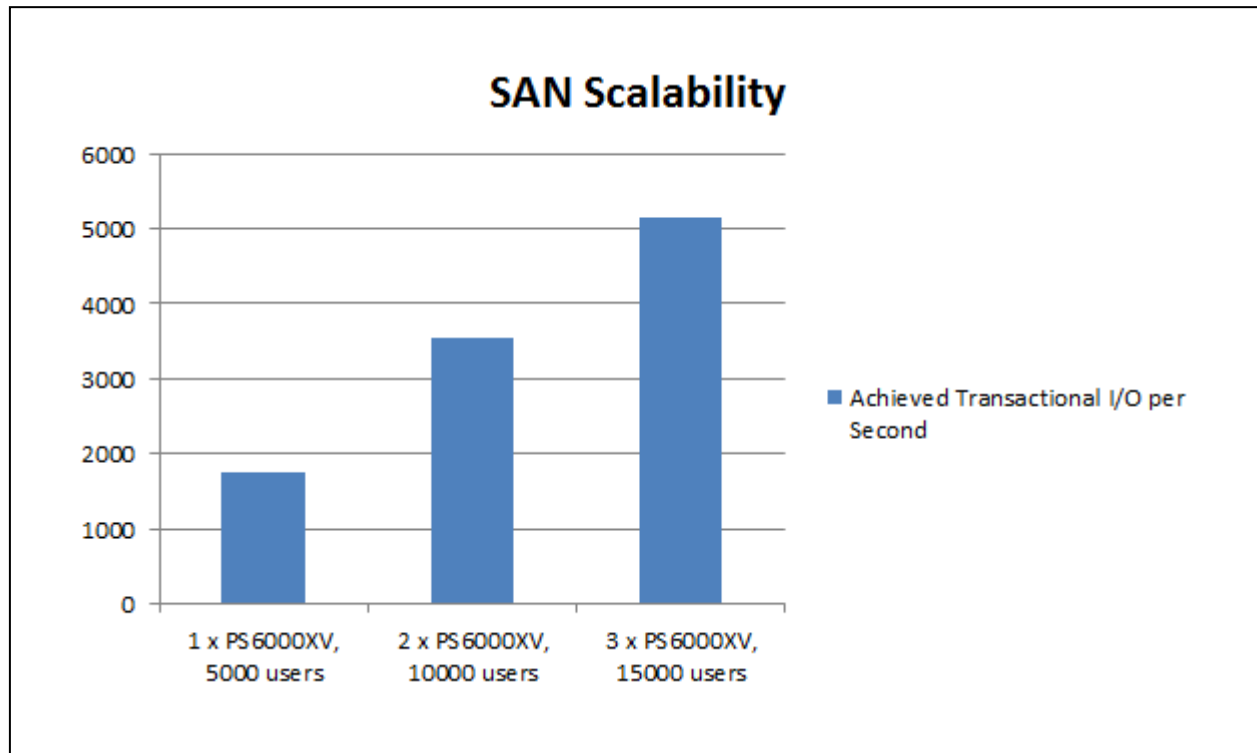


Figure 11 SAN Scalability Study - Average Database IOPS

member test configurations. This is because the EqualLogic Peer Storage architecture scales the available storage resources linearly to provide the required IOPS. The storage processing resources in EqualLogic storage architecture that linearly scale as we add more arrays include: the number of disk drives, the number storage processing controllers including on-board cache, and the number storage network interfaces. The number of network interfaces on the server side was kept constant and this did not matter since the workload was random I/O and therefore was more disk bound and storage processor bound. The network switches were capable of supporting the arrays under test. If additional arrays are deployed for scalability, a different network switch configuration may be required based on port count requirements and workload characteristics.

3 Sizing EqualLogic Storage for Exchange Server 2010 Mailbox Role in a Virtualized Environment

In an Exchange 2010 environment, storage sizing will depend on capacity and performance (IOPS) requirements of the mailbox servers. These requirements are determined mainly by the number of users, the mailbox sizes and the mailbox IOPS profiles that need to be met. Other operational factors within Exchange Server 2010 systems described later in this section can also affect sizing requirements. The primary storage components that need to be sized correctly are the databases and the logs. The results from I/O Mailbox Profile studies in Section 2 were used for determining the sizing guidance in this section.

3.1 Exchange Random IOPS and EqualLogic Storage

Exchange 2010 uses 32KB block I/O and the I/O pattern is random. The acceptable latency limits for Exchange Server database read IOPS is 20ms and for Log write IOPS is 10ms⁹.

The maximum number of random IOPS that an EqualLogic PS Series storage array can provide will depend on the disk drive rotational speeds, the number of drives in the array and the RAID type.

3.1.1 Disk Drive Rotational Speed

As shown in Section 2.3, higher RPM drives can provide higher array IOPS. Among the three drive types offered in PS Series arrays, SAS 15K drives offer the highest level of random I/O performance measurable as IOPS. The EqualLogic PS6000XV array supports 15K RPM SAS drives. SAS 10K drives offer lower level of random I/O performance than SAS 15K drives. EqualLogic PS6000X and PS6500X arrays support 10K RPM drives. SATA II drives are the lowest performing drives offered in EqualLogic PS Series arrays. The PS6000E and PS6500E arrays support 7.2K SATA II drives.

If your IT department has standardized on a particular type of drive for deployments, then the choice would be pre-determined. If there is no preference for a particular drive type, then our recommendation would be to use the sizing process guidelines in this document.

3.1.2 Number of Disk Drives in the Array

Another factor that can affect array IOPS is the number of drives in the array. EqualLogic PS Series storage arrays come in two drive count form factors: the PS65x0 series (supports 48 drives per enclosure) and the PS60x0 series (supports 16 drives per enclosure).

3.1.3 RAID Policy

PS arrays support various RAID levels that include RAID 5, RAID 50, RAID 10 and RAID 6. RAID 50 offers a good combination of useable capacity and performance as compared with RAID 5 or RAID 6. In RAID 50, application writes are balanced across multiple RAID 5 sets to reduce write overhead performance impact. Hence RAID 50 is recommended as compared with RAID 5.

RAID 10 provides more IOPS than RAID 50 for the same type of workload. This is because RAID 50 incurs parity calculation along with necessary reads and writes to disk (4 disk I/O operations) to

⁹ Temporary spikes up to a maximum of 100ms are considered acceptable. Reference: [http://technet.microsoft.com/en-us/library/ff367871\(EXCHG.141\).aspx](http://technet.microsoft.com/en-us/library/ff367871(EXCHG.141).aspx)

support a single application write operation. In comparison RAID 10 just mirrors or writes an application write operation to a second disk drive (2 disk I/O operations per application write operation).

Table 11 shows measured Exchange 2010 database IOPS levels for EqualLogic PS Series arrays that will meet the maximum latency requirements¹⁰ for Exchange Server 2010 database I/O deployed in a VMware vSphere virtualized environment.

Note: The values in Table 11 are based on the results of the Jetstress tests run on the specific system configurations described in Section 2. In other environments, measured IOPS will vary based on the actual configuration deployed. These values should be used as a starting point for validating results of further testing using Exchange parameters (mailbox profile characteristics, database layout, system resources) that match your environment.

Measured Transactional IOPS ^(a)		
Model ^(b)	RAID 50	RAID 10
PS6000XV	1745	2165
PS6000X	1100	1365
PS6500E	1800	2400

(a) Transactional IOPS in this table are based on tests run with background database maintenance (up to 4 databases). The values in this table correspond to use of a single EqualLogic PS Series array which is dedicated to hosting Exchange data only. If other applications in the virtual environment share the same array(s) for hosting data, then these values should not be used.

(b) SSD drives supported by the EqualLogic PS6000S are not considered due to the reduced IOPS requirement with Exchange Server 2010 as compared to earlier versions, plus cost economics and lower SSD drive capacity levels.

Table 11 Recommended Exchange Database IOPS Based on RAID Policy and Array Type

Other RAID policy factors that you should consider:

- Degraded mode operation due to drive failure: RAID 10 provides better response times than RAID 50 in degraded situations when one of the drives in a RAID set fails.
- RAID 10 sets complete rebuilds faster than RAID 50 sets when a failed disk is replaced with a new one.

3.2 EqualLogic Storage – Usable Capacity

Usable capacity available from a storage array will depend on the drive size within array, drive count and RAID type primarily. Generally SATA II drives offer higher levels of capacity per drive as compared to SAS 15K or 10K RPM drives. The PS6500 series supports a total of 48 drives as compared to PS6000 series with 16 drives.

¹⁰ The acceptable latency limits for Exchange Server database read IOPS is 20ms and for Log write IOPS is 10ms.

RAID 50 offers more useable space than RAID 10, since half the number of disk drives in RAID 10 set is used for mirroring. On the contrary, only a few disk drives are used in RAID 50 for storing parity information. Table 12 shows the raw space available on each array according to disk drive sizes and PS Series model type.

Model	Disk Drive Size & (Model)	RAID 50 Useable Capacity (GB)	RAID 10 Useable Capacity (GB)
PS6000 Series	250GB (E)	3000	1750
	300GB (X , XV)	3600	2100
	400GB (X)	4800	2800
	450GB (XV)	5400	3150
	500GB (E)	6000	3500
	600GB (X, XV)	7200	4200
	1TB (E)	12000	7000
	2TB (E)	24000	14000
PS6500 Series	500GB (E)	19500	11500
	600GB (X)	23400	13800
	1TB (E)	39000	23000
	2TB (E)	78000	46000

Table 12 PS Series Array Capacity¹¹

3.3 Sizing EqualLogic Storage for Microsoft Exchange

In this section we discuss general guidelines for Exchange storage capacity sizing. We also present sizing processes and formulas that can be used to calculate storage capacity and IOPS requirements.

3.3.1 Microsoft Guidelines for Storage Capacity Sizing

Mailbox server database capacity will depend primarily on mailbox quota size, database white space and mailbox database recoverable items. Together, these factors along with storage quota and associated overheads determine the actual mailbox size for a user. Additional operational factors impacting database size are space for content indexing, offline database maintenance space, recovery database needs and overhead space for future growth or contingency. Backup SLA time requirements and space for database copies to support high availability will also impact capacity calculations. Mailbox server log capacity primarily depends on the user messaging activity profile (messages sent/received per day).

¹¹ Useable capacity will be slightly less than the values in this table due to space requirements for PS Series array management functionality, and file system/operating system requirements.

Each transaction log file in Exchange 2010 has a size of 1MB. Per Microsoft® guidelines, the number of log files increase by a factor 1.9 if the average message size doubles to 150KB and by factor 3.8 if message size is above 150KB. Other log capacity factors include:

- Log retention requirements to support backup schedules
- Move mailbox operations
- Log overhead for future growth or contingency.

Additional factors¹² per mailbox server include:

- Additional log capacity for database copies
- Replay lag space for copies
- Log truncation mechanisms

Note: Microsoft® provides detailed capacity sizing guidelines for databases and logs in the following TechNet article:

- *Understanding Mailbox Database and Log Capacity Factors:*
<http://technet.microsoft.com/en-us/library/ee832796.aspx>

3.3.2 Microsoft Guidelines for Storage Performance Sizing

Database storage performance sizing depends on the messaging profile. The messaging profile is defined as the average number of messages sent/received by each user per day. The messaging profile determines the number of IOPS generated by each user. Estimated IOPS values also depend on the database cache and memory allocated per user. Microsoft® provides tables for estimating database cache and IOPS per mailbox based on messaging activity. The guidelines¹³ provided by Microsoft® are for users running in Outlook cached mode and include database and content indexing IOPS.

Note: Microsoft® provides a detailed discussion of database and log performance factors, including the IOPS generated for a range of messaging profiles, in the following TechNet article:

- *Understanding Database and Log Performance Factors:*
<http://technet.microsoft.com/en-us/library/ee832791.aspx>

Additional IOPS estimation factors can include:

- Desktop search engines on client machines
- Third-party applications including mobile messaging, client outlook mode (online / cached)
- Content indexing I/O for search cataloging.

Non-transactional (non-user related) I/O impacts on the database can include:

¹² More details are available in the following Microsoft® TechNet article: <http://technet.microsoft.com/en-us/library/ee832796.aspx>

¹³ See <http://technet.microsoft.com/en-us/library/ee832791.aspx>

- I/O for background maintenance
- Online maintenance and message records management (MRM) activities.

Overhead IOPS for contingency or future growth need to be considered as well (20% recommended). If database copy is enabled, then log write I/O can typically comprise 40% to 50% of database write IOPS in Exchange 2010. Additional log I/O will be incurred for database copies for logs to be read and replicated to copy servers. This is typically 10% of active log writes per database copy. Overhead IOPS (typically 20%) need to be considered as well.

3.3.3 Database Capacity Sizing Process

The goal of this process is to calculate an estimated total database size. Table 13 shows the recommended process to follow.

Variable(units)	Formula	Definition/notes
Step #1: determine user profile variables:		
NU		Number of mailbox users (including future growth) ¹⁴
MQ (MB)		Mailbox storage Quota or limit, per user
MSR		Number of Messages Sent/Received per day (the mailbox profile)
AMS (KB)		Average Message Size
Step #2: calculate White Space per User (WSU) :		
WSU (MB)	= MSR x AMS/1024	White space per user.
Step #3: Calculate Dumpster Size per User (DSU) :		
DSU (MB)	= MSR * AMS/1024 * 14 + (0.012 * MQ) + (0.058 * MQ)	<ul style="list-style-type: none"> • By default, Exchange 2010 stores deleted items for 14 days. • Single item recovery requires 1.2% additional space¹⁵. • Calendar version logging (enabled by default) adds 5.8% additional space.
Step #4: Calculate Mailbox size per User (MBU) :		
MBU (MB)	= MQ + WSU + DSU	Actual mailbox size per user
Step #5: Calculate estimate for Total Database Size (DBS):		
DBS (MB)	= NU * MBU * 1.2	Projected total database size. We recommend using a 20% overhead factor.
Step #6: Calculate estimate for Individual Database Size (IDBS) :		
IDBS (MB)	= DBS / DB Count	See the section "Estimating Database Count" following this table.
Step #7: Calculate estimate for Database Volume Size (DVS) :		

¹⁴ Excludes mobile user requirements

¹⁵ For more details, see Microsoft® TechNet Article "Understanding Mailbox Database and Log Capacity Factors": <http://technet.microsoft.com/en-us/library/ee832796.aspx>

DVS (MB)	= IDBS * 1.1 * 1.2	<ul style="list-style-type: none"> • Content indexing requires an additional 10% of space per database within the same volume. • After all factors are considered 20% overhead space within the database volume is recommended.
<p>Additional sizing factors to consider:</p> <ul style="list-style-type: none"> • Offline database maintenance, if required, will require an additional 110% of volume space. • Recovery databases, if planned, will need additional space based on the number of databases to be recovered in the separate space. 		

Table 13 Calculation Process - Total Database Size Estimation

Estimating Database Count

Before you can calculate IDBS, you need to determine the database count. The number of databases is typically determined based on backup requirements. The size of a single database can be determined based on backup needs such as backup throughput and SLAs. If you are not performing full daily backups or using Exchange Server 2010 DAG (Database Availability Groups) for high availability and restore, then the requirements calculator¹⁶ provided by Microsoft® can be used to determine database count.

In a DAG model, Microsoft® recommends database sizes up to 2TB. In standalone mode, the database size recommendation is 100GB. Sometimes, instead of size, the maximum number of users to be managed per database may determine the database count. Also, the increased manageability of multiple database configurations could be a factor in determining database count. The particular edition (Standard or Enterprise) of Exchange server 2010 will limit the database count as well. If 24 x 7 background database maintenance is enabled, additional I/O will be incurred per database deployed. If the number of databases is high, then the background database maintenance I/O will have an impact on normal mailbox I/O. With high database counts, this becomes an important consideration.

Recommendation: After determining database count, deploy one database per volume in EqualLogic Storage. This enables isolation of management and backup operations for each database to a single volume. In this case, the number of storage volumes required is equal to the database count determined above.

Additional optional factors: Offline database maintenance, if required, will need 110% of target volume space. Recovery databases, if planned, will need additional space based on the number of databases to be recovered in separate space.

3.3.4 Log Capacity Sizing Process

The goal of this process is to calculate an estimated Log Volume Size. Table 14 shows the recommended process to follow. Additional considerations for the calculation process shown in Table 14:

- Different backup models such as weekly full backups would warrant log retention for more than 3 days. In that case, more space needs to be provisioned based on the backup period.

¹⁶ See <http://msexchange.com/files/12/attachments/entry453145.aspx>

- In a DAG model, the database copy can have a lag in replicated log content from the source database specified by number of lagged days (the replay lag time). Based on this, more space will need to be provisioned for logs depending on the replay lag period.
- Move mailbox operations will require more log space based on the number of mailboxes to be moved across databases periodically for maintenance operations.

Recommendation: After determining database count, you should deploy one log volume per database in EqualLogic Storage. This enables isolation of database and logs for management, recoverability and backup operations via hardware snapshot providers such as the EqualLogic Auto-Snapshot Manager (ASM) for Microsoft® Exchange. However, this approach may require more drive letter assignments to mount all volumes, and the use of volume mount points may be necessary. See <http://technet.microsoft.com/en-us/library/ee832794.aspx> for more information on Exchange LUN architecture.

Variable(units)	Formula	Definition/notes
Step #1: Define process inputs:		
LFS (MB)	= 1.0 (constant)	Exchange 2010 log file size is 1MB.
DB Count	(from database capacity sizing process)	See the section "Estimating Database Count"
NU		Number of mailbox Users (including future growth)
MSR		Number of Messages Sent/Received per day (the mailbox profile)
AMS (KB)		Average Message Size
Step #2: Define Number of transaction Logs generate per User per day (NL):		
NL	= MSR * 0.2	
Step #3: Calculate Message Size Factor (MSF) :		
MSF	= 1.0 (AMS <= 75KB) = 1.9 (AMS <= 150KB) = 3.8 (AMS > 150KB)	
Step #4: Calculate Number of users per Database (NDB):		
NDB	= NU / DB Count	
Step #5: Calculate Total Log Capacity per day per Database (TLC):		
TLC (MB)	= NDB * NL * MSF * LFS	
Step #6: Calculate estimate for Log Volume Size (LVS):		
LVS (MB)	= TLC * 3 * 1.2	<ul style="list-style-type: none"> • Total log capacity per volume will affected by backup and high availability options. • If daily backups are taken, then it is a best practice to provision space for 3 days of logs. The same recommendation applies if a DAG deployment model is used. • We recommend using an additional 20% overhead factor.

Additional sizing factors to consider:

- Offline database maintenance, if required, will require an additional 110% of volume space.
- Recovery databases, if planned, will need additional space based on the number of databases to be recovered in the separate space.

Table 14 Calculation Process - Log Volume Size Estimation

3.3.5 Database I/O Performance Sizing Process

The goal of this process is to calculate an estimated total database IOPS requirement. Table 15 shows the recommended process to follow.

Variable(units)	Formula	Definition/notes
Step #1: Define process inputs:		
MSR	(from database capacity sizing process)	Number of Messages Sent/Received per day (eg. the mailbox profile)
NU		Number of mailbox Users (including future growth)
Step #2: Define Standalone Database Overhead (if not deployed in a DAG) (SDBO):		
SDBO	= 1.2 (if DAG not used) = 1.0 (if DAG used)	Derived from recommendations provided by Microsoft® TechNet ¹⁷ .
Step #3: Calculate IOPS per User (IOPU) :		
IOPU	= SDBO * (MSR / 1000)	
Step #4: Calculate estimate total Database IOPS required (DBIOP) :		
DBIOP	= NU * IOPU * 1.2	A 20% overhead is recommended for sizing IOPS.

Table 15 Calculation Process - Database I/O Performance Sizing

3.3.6 Log I/O Performance Sizing Process

The goal of this process is to calculate an estimated total database IOPS requirement. Table 16 shows the recommended process to follow.

Variable(units)	Formula	Definition/notes
Step #1: Define process inputs:		
DBIOP	(from database I/O performance sizing process)	Total database IOPS required
Step #2: Define Total Log IOPS (LIOPS) :		
LIOPS	= (DBIOP * 0.5 * 0.5) + (DBIOP * (number of DAG copies - 1) * 0.1) * 1.2	

¹⁷ <http://technet.microsoft.com/en-us/library/ee832796.aspx>

LIOPS formula considerations:

- The amount of log I/O is about 50% of the database write I/O.
- Database write I/O is about 40% to 60% of total database I/O based on number of messages per user and database high availability model. An average value of 50% will fit most cases.
- An additional 10% log I/O will be incurred for each additional database copy in a DAG
- A 20% overhead factor is recommended.

Table 16 Log I/O Performance Sizing Process

3.4 Choosing EqualLogic Array

Follow these steps to choose an EqualLogic array for Exchange Server 2010 deployments. These steps assume a DAG scenario with at least 2 database copies (primary and secondary).

According to Microsoft® recommended best practices for Exchange 2010, it is okay to host databases and logs on the same RAID set and drives in a DAG scenario. However, copies of a database must be hosted on separate drives or RAID sets.

The following selection process will recommend the same approach, as follows:

- Database and log volumes are hosted on same array(s)
- Two database copies in a DAG deployment model (primary and secondary databases) spread across separate arrays.

3.4.1 Selection Process

1. Determine the RAID type needed (RAID 50 or 10) based on recommendations in Section 3.1.
2. Calculate the total Exchange Server capacity needed using the total number of database and log volumes needed and the size of each database and log volume. Other storage requirements besides Exchange server capacity requirements need to be considered as well. For example, space required for volume snapshots reserve or volume replication to support backup and restore operations. Also, any space required for cloned volumes would be an additional factor.
3. Based on total capacity needed, use Table 12 to determine type and quantity of storage array(s) required. Choose the appropriate drive size within an array based on useable capacity from the array and total capacity needed.
4. Determine if the number of arrays of the same model chosen in step 3 will satisfy total database IOPS (DBIOP) needed based on the values in Table 11. Consider any additional IOPS needed in the DBIOP value – for example IOPS for mobile users. You may need to add additional arrays of the same model to satisfy the user IOPS needs. Table 11 includes IOPS for logs within the same array as well as database maintenance (for up to four databases). Do not add log IOPS to the total requirement when using this table.
5. If step #4 did not provide satisfactory solution or the array model chosen in step #3 does not match the IOPS needs in step #4, go back to step #3 and choose a different array model.
6. Once an appropriate array model with correct drive size is determined, and the total number of such arrays is determined, you will need to deploy twice the number of arrays if you require support for database copies. If more than two database copies are required, scale up the number of arrays needed as necessary.

In standalone mode, it is a generally accepted best practice (for recoverability purposes) to have databases and logs deployed on different RAID sets and physical drives. . However this may not be possible in smaller deployments with limited storage resources. In smaller environments using Dell EqualLogic storage arrays, the transaction log volumes can be hosted within the same pool and disk drives as the database volumes without performance loss. EqualLogic arrays provide high reliability via RAID and redundant hot swappable components. Additional data backup solutions for offline data copies can be implemented for supplemental protection. However, if high levels of data protection or recoverability are required, as is usually the case in large environments, data replication via Exchange DAGs should be used.

3.4.2 Sizing Example for an Exchange 2010 Environment Supporting 5000 users

In this example we will use the following User Profile:

- 5000 users
- 2GB mailbox quota size
- 150 messages sent/received per day per user (very heavy user profile)
- Average message size = 75KB

We will use the following system parameters:

- The system will be designed for high availability, using Database Aggregation Groups (DAG): 2 copies
- Maximum database size = 1024GB

Using the Database Capacity Sizing Process in 3.3.3:

NU	= 5000	
MQ	= 2048MB	
MSR	= 150	
AMS	= 75KB	
WSU	= 11MB	White space per user: 150 x 75/1024
DSU	= 297MB	Dumpster size per user: (150 x 75 x 14 / 1024) + (2048 x 0.012) + (2048 x 0.058)
MBU	= 2356MB	Mailbox size per user: 2048MB + 11MB + 297MB
DBS	= 11505GB	Total database size: 5000 x 2356
DB	= 12	Database Count: (DBS / max DB size: 11505 / 1024 = ~12)
IDBS	= 959GB	Individual database size: 11505 / 12
DVS	= 1266GB	Database volume size: 959 x 1.1 x 1.2

Table 17 Example: Database Capacity Sizing Process

Using the Log Capacity Size Estimation process in Section 3.3.4:

DB Count	= 12	Therefore number of log volumes = 12
NU	= 5000	
MSR	= 150	
AMS	= 75KB	
NL	= 30	Number of transaction logs generated per user per day: $150 * 0.2 = 30$
MSF	= 1	Message size factor: 1 (AMS <= 75KB)
NDB	= 417	Number of users per database: $5000 / 12$
TLC	= 12510MB	Total log capacity per day per database: $30 * 417 * 1 * 1MB$
LVS	= 48GB	Log volume size: $12510 * 1.1 * 3 * 1.2 / 1024$ (TLC x 10% for mailbox move x 3 days x 20 % overhead) <ul style="list-style-type: none"> • Assume maximum 10% of users moved within log truncation window for maintenance operations. • Recommended log truncation failure protection (due to backup or replication) is 3 days and overhead space is 20%.

Table 18 Log Volume Capacity Sizing Process

For 12 database volumes and 12 log volumes, the total capacity needed for all volumes can be calculated using values from Table 17 and Table 18 as follows:

Total capacity needed for all volumes = 15768GB (database volume size x 12 + log volume size x 12)

To estimate total IOPS required we estimate IOPS required per user = 0.15. Based on that estimate, we calculate total database IOPS required:

Total database IOPS required = $5000 * 0.15 * 1.2 = 900$ IOPS (number of users x IOPS per user x 20% overhead)

Based on a total capacity requirement of 15780GB (from Table 12), the PS6500 array with 500GB SATA drives in a RAID 50 configuration can be chosen. Based on a total database IOPS requirement of 900 (from Table 11), the PS6500 can satisfy the IOPS requirement of 900 IOPS for the user profile under consideration. Hence the PS6500 can be a good fit for this profile. Two PS6500 will be needed to support an active and passive copy for each database. The active copy of a database should be hosted on one array and the passive on the other array.

Processor and memory sizing recommendations are provided by Microsoft®, in the following TechNet article:

- **Mailbox Server Processor Capacity Planning:** <http://technet.microsoft.com/en-us/library/ee712771.aspx>

4 Exchange 2010 Database Availability Group (DAG) Studies

In the second phase of our lab tests we conducted a series of I/O profiling studies to characterize mailbox server I/O and Database Availability Group (DAG) replication I/O in an Exchange 2010 environment.

4.1 Workload Simulation and Server Performance Criteria

We used Microsoft® Exchange Load Generator 2010¹⁸ (Loadgen) to simulate a client I/O workload in the Exchange test configuration. In all test cases, Loadgen was run with the following loading attributes:

- We simulated a workload of 5000 simultaneous Outlook 2007 users¹⁹ in online mode, with search operations turned off, for a duration of four hours.
- Background database maintenance and scheduled maintenance was turned off.

Our Exchange server performance criteria for the Loadgen testing consisted of the following:

- Average disk latency for active database volumes maintained less than 20ms, with maximum measured latency less than 100ms.
- Mailbox server counter "MSExchangeIS\RPC Average latency" less than 10ms on average.
- Mailbox server counter "MSExchangeIS\RPC Requests" less than 70 at all times.
- Mailbox server counter "MSExchangeIS Mailbox(_Total)\Messages Queued For Submission" less than 50 at all times.

4.2 Test Configuration

The test configuration for the Loadgen studies is shown in Figure 12 and Figure 13 below. The actual server connectivity to storage is identical to the test configuration shown in the Jetstress studies (Section 2.1).

- Two PowerEdge M710 servers were deployed for Exchange Server 2010 roles
- Each M710 server hosted two virtual machines – one for the Exchange Server 2010 mailbox role and another for Exchange Server 2010 Hub Transport / Client Access Server (CAS) Roles.
- The Mailbox Server VM was configured with 8 vCPUs and 48 of memory (reserved in ESX host). The Hub / CAS VM was configured with 4 vCPUs and 6GB of memory (reserved). Storage connectivity is similar to the illustration in Jetstress studies above.
- The two mailbox server VMs across the two blade servers were configured as members of an Exchange 2010 DAG (Figure 12 below). If one of the servers incurs failure, the other server in the DAG can continue hosting databases that were active on the failed server via a fail-over process. Passive database copies are maintained on the second server and kept synchronized

¹⁸ See Microsoft® TechNet Article: "*Tools for Performance and Scalability Evaluation*", <http://technet.microsoft.com/en-us/library/dd335108.aspx>

¹⁹ Using the Loadgen "Exchange 2007 Online Very Heavy" user profile

with the active copies via log replication. One of the Hub/CAS servers was utilized for configuring the file share witness for the DAG.

- The two Hub / CAS VMs were configured as members of a Network Load Balancing (NLB) cluster (unicast mode and single affinity). One virtual NIC was used for NLB and another for public access. NLB was used to load balance the Exchange CAS services (Figure 13 below)
- Each Hub/CAS VM was configured with two virtual NICs for server LAN connectivity – one for public access and other for NLB private connectivity (unicast NLB mode). Also each VM was configured with 4 virtual NICs of type vmxnet3 for iSCSI connectivity. The Hub Transport queue database and logs were hosted on a separate volume on the same group member array used by their respective mailbox servers.
- Each mailbox VM was configured with two virtual NICs for server LAN connectivity – one for public access and other for DAG replication. Also each mailbox VM was configured with 4 additional virtual NICs of type vmxnet3 for iSCSI connectivity using the iSCSI initiator in the VM.
- The Microsoft® iSCSI initiator in the guest Windows Server 2008 R2 OS was used for storage connectivity. The VMs were installed with EqualLogic MPIO DSM (HIT kit) for EqualLogic aware multi-path I/O to the EqualLogic SAN.
- On each ESX host, the server-side LAN virtual NICs for all VMs were connected to vSwitch0 and storage iSCSI virtual NICs of all VMs were connected to vSwitch1. vSwitch0 used the onboard Broadcom 5709 1GbE NICs on server Fabric A to connect to the external M6220 switches for server LAN access. vSwitch1 used the 1GbE Broadcom 5709 NICs on fabric B to connect to the external M6348 switches for storage access.

Two PowerEdge M610 servers on the same blade M1000e chassis were deployed for other test components. One M610 server hosted the Active Directory VM and the vSphere Virtual Center VM. The second M610 server hosted the client VM used for executing the simulation tool Loadgen and the monitoring VM for running Windows PerfMon and EqualLogic SAN Headquarters.

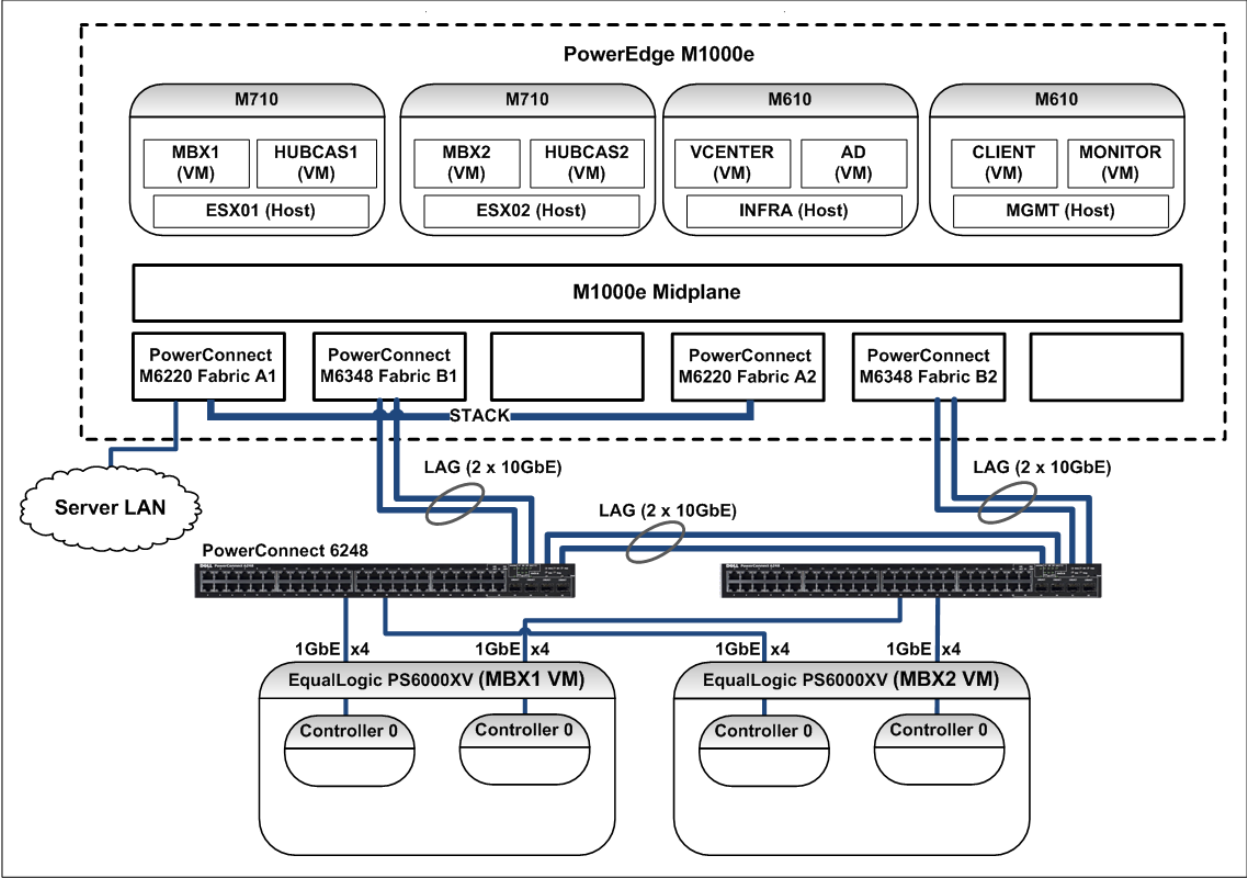


Figure 12 Test Configuration: DAG Loadgen

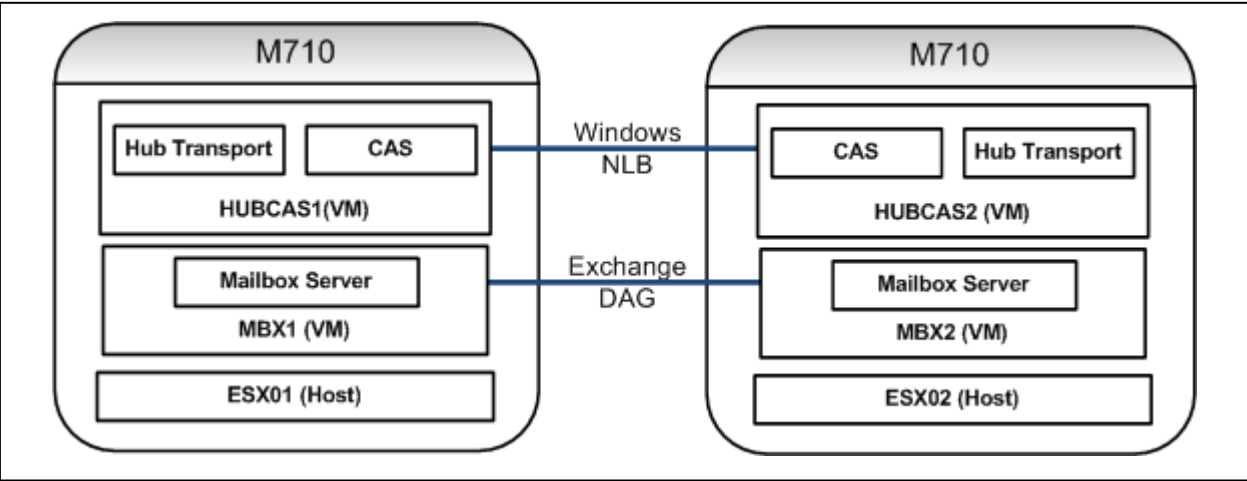


Figure 13 Test Configuration: NLB Cluster Detail

4.3 Active/Passive Configuration

In the first test configuration, all active databases in the DAG were hosted on one Mailbox Server and its corresponding PS Series member within the EqualLogic SAN group. The second mailbox server hosted the passive database copies on a second group member. A single EqualLogic PS Group was configured with two storage pools, Pool A and Pool B, with each pool comprised of one storage array. Each of the two mailbox servers in the configuration connected to database and log volumes in respective pools on the EqualLogic SAN. Since Exchange is configured in copy mode, the database and log volumes can be hosted on the same RAID set (EqualLogic member). Four mailbox databases were created with one database and one log volume each within one storage pool. The four database copies also included a database and log volume each on the second storage pool. This database and storage configuration is shown in Figure 14.

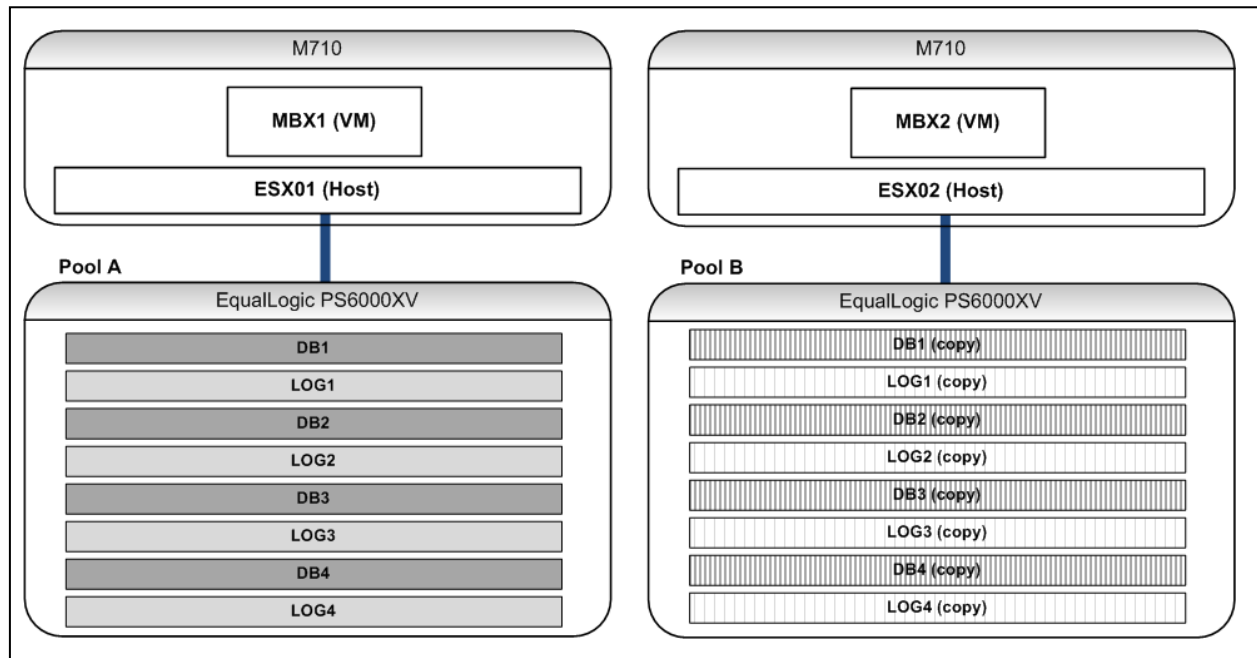


Figure 14 DAG Active/Passive Volume Configuration

4.3.1 Active/Passive Configuration Test Results

Figure 15 shows virtual machine host (ESX) processor and memory utilization levels measured using the VMware `esxtop` utility. ESX01 is the VMware ESX server hosting the active MBX1 VM. ESX02 is the VMware ESX server hosting the passive MBX2 VM. ESX01 in this test showed slightly higher CPU utilization due to its role as the active mailbox server.




Figure 15 DAG Active/Passive Resource Utilization

Figure 16 shows the database IOPS levels for the active mailbox server (MBX1) during the test. In Figure 16 the “Attached” values correspond to the active database, and the “Recovery” values correspond to the passive database copies. (Since there were no passive database copies hosted on MBX1, the IOPS values for the “Recovery” components were zero.) The total average mailbox server IOPS on MBX1 during the test duration was about 449 IOPS (sum of each of the average components in Figure 16). This I/O load was generated to satisfy end user (Outlook client) operations.

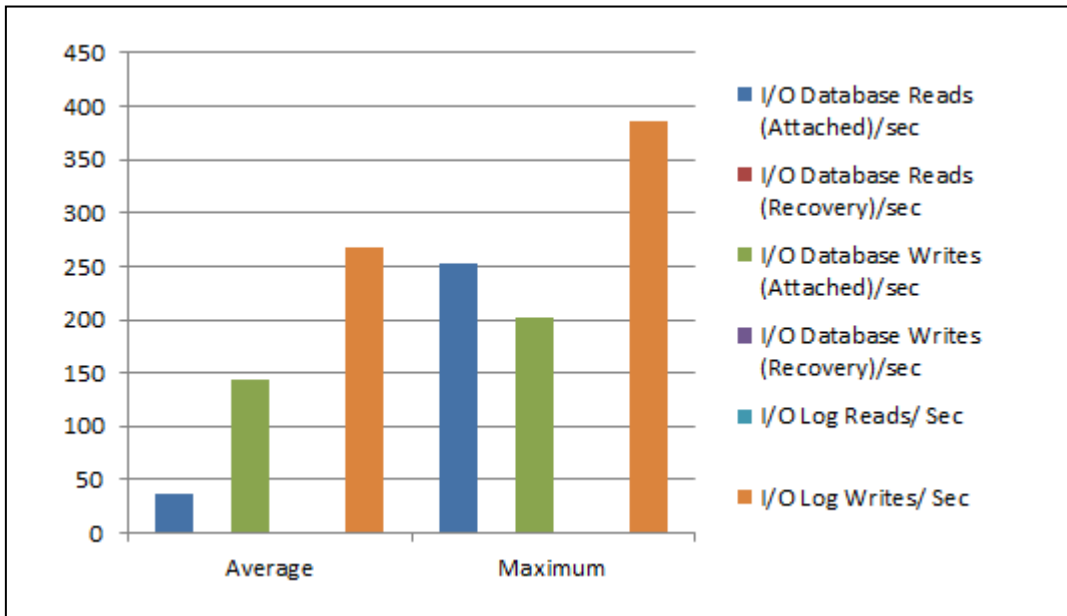


Figure 16 DAG Active Database I/O (Active/Passive Configuration)

Figure 17 shows the database IOPS levels for the passive mailbox server (MBX2) during the same test. Since there were no active database copies, the “attached” components in the chart are zero. The total average mailbox server IOPS on MBX2 during the test duration was approximately 322 IOPS (sum of each of the average components in Figure 17 below). This I/O was generated to keep the passive database copies synchronized with the active databases. Logs generated on the active copies were copied by the Exchange Server replication component to the passive server and then replayed on the passive copies.

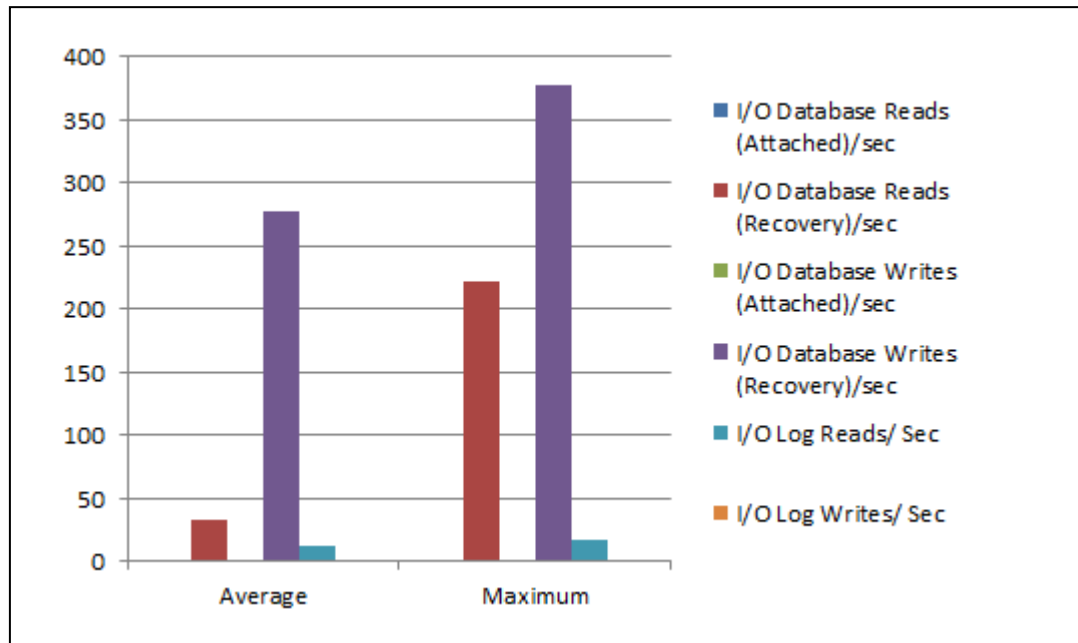


Figure 17 DAG Passive Database I/O (Active/Passive Configuration)

4.4 Active/Active Configuration

In the Active/Active test configuration we split the active databases in the DAG between the two Mailbox Servers such that both storage pools A and B hosted two active databases and two database copies. This configuration is shown in Figure 18.

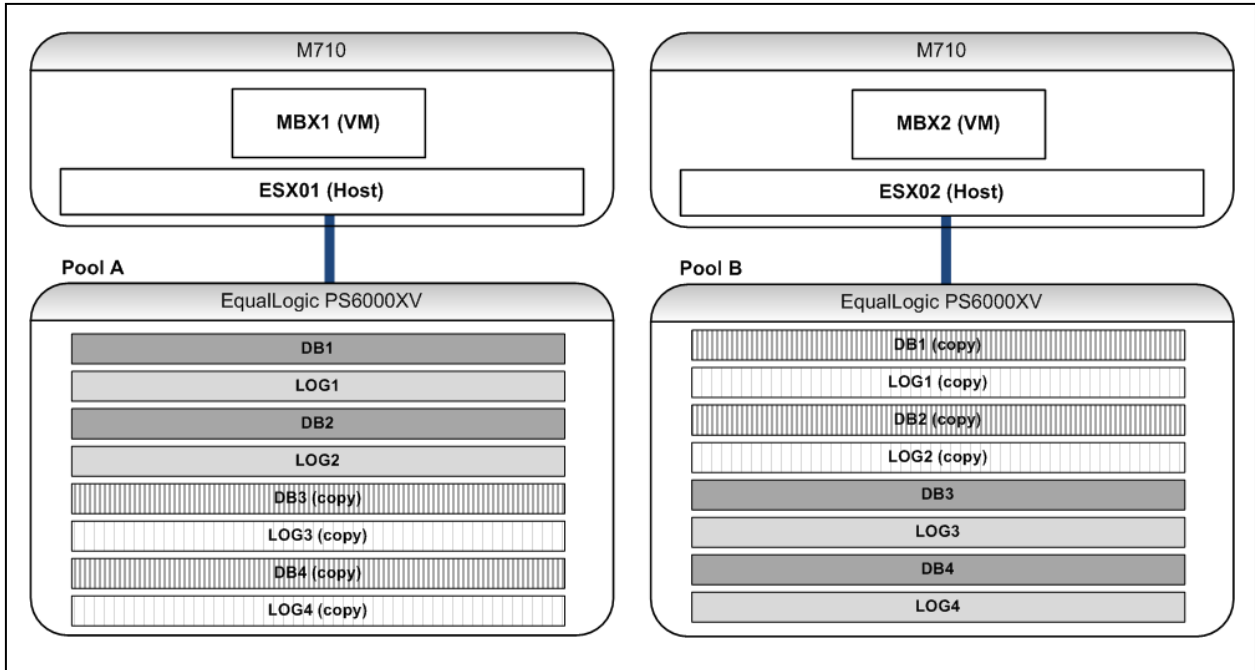


Figure 18 DAC Active/Active Volume Configuration

4.4.1 Active/Active Configuration Test Results

Figure 19 shows virtual machine host (ESX) processor and memory utilization levels measured using the VMware esxtop utility. ESX01 is the VMware ESX server hosting the active MBX1 VM. ESX02 is the VMware ESX server hosting the passive MBX2 VM.

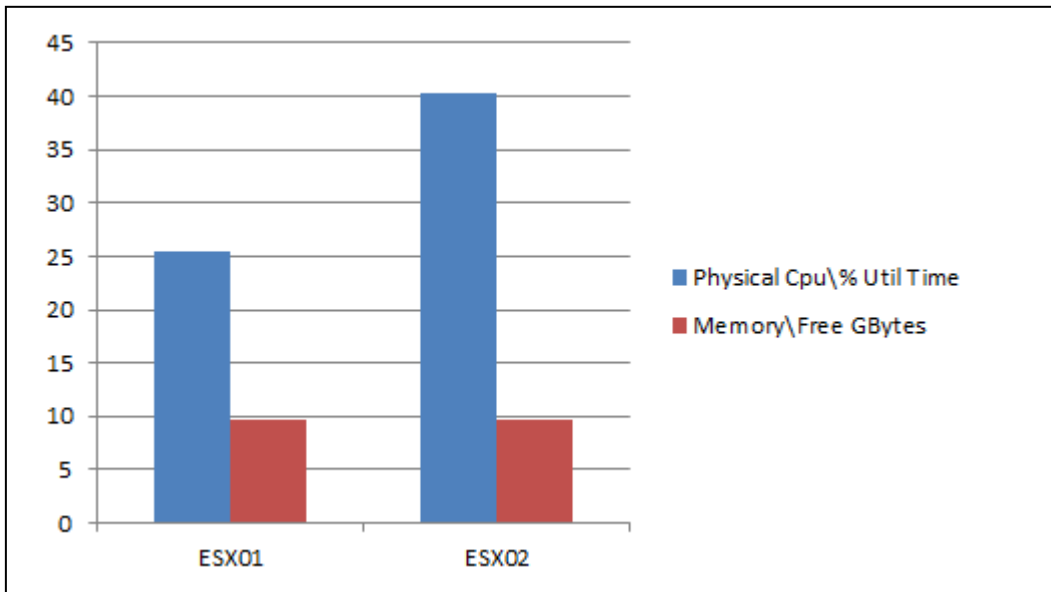


Figure 19 DAG Active/Active Resource Utilization

The processor utilization for ESX02 shown in Figure 19 was caused by the extra VM processor load generated by the HUBCAS2 VM (see Figure 13). The higher load on HUBCAS2 was caused by the single Sizing and Best Practices for Deploying Microsoft Exchange Server 2010 on VMware vSphere and Dell EqualLogic Storage

affinity option setting in windows NLB services. This setting directed all incoming requests from the Loadgen client to the HUBCAS2 VM.

The database IOPS for mailbox servers MBX1 and MBX2 are shown in Figure 20 and Figure 21.

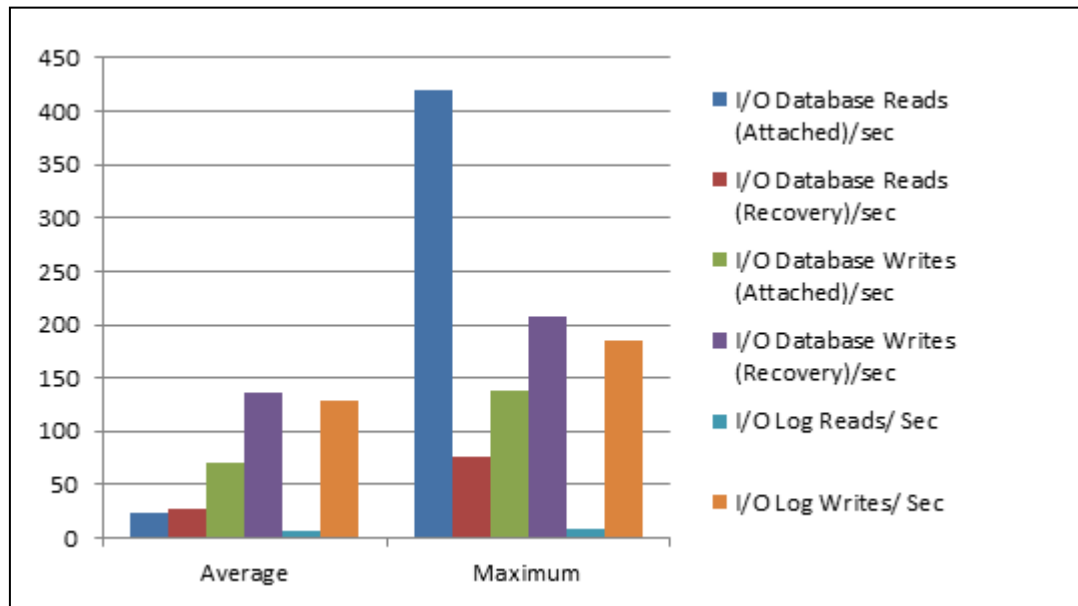


Figure 20 DAB MBX1 Database I/O (Active/Active Configuration)

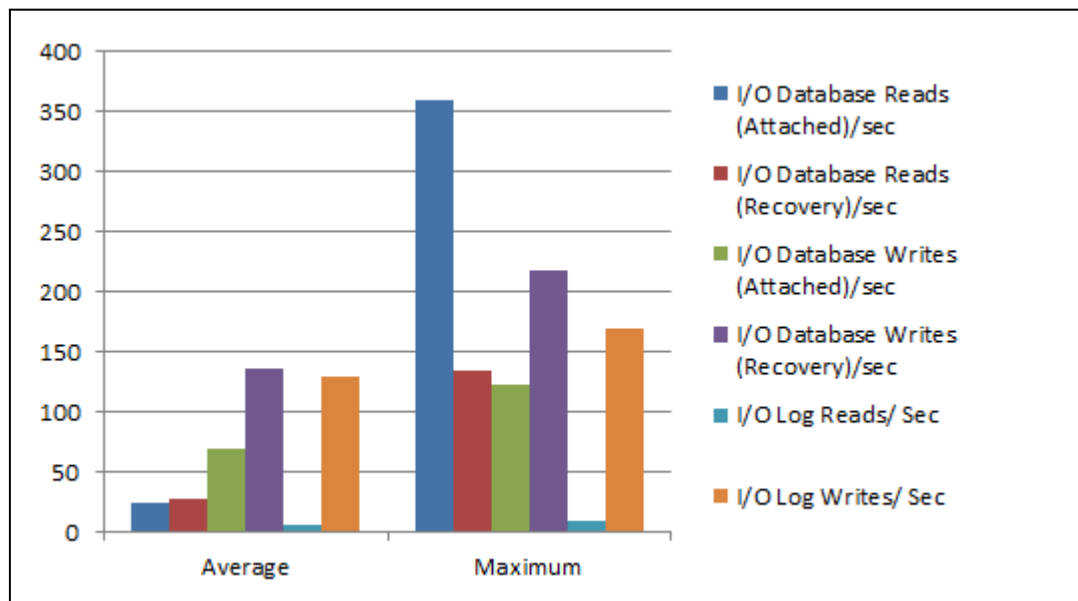


Figure 21 DAB MBX2 Database I/O (Active/Active Configuration)

Figure 20 and Figure 21 show that the measured IOPS profiles for both mailbox servers in the Active/Active configuration are virtually identical. The results show both servers processing an equivalent amount of IOPS for both active (“attached”) and passive (“recovery”) databases. The

variation in total average mailbox server IOPS between MBX1 and MBX2 was less than 1% (390 IOPS for MBX1 and 392 IOPS for MBX2).

As a best practice we recommended deploying Exchange in an Active/Active configuration similar to that shown in Figure 18. In this case it is clear that the Active/Active configuration enabled more efficient utilization of system resources.

Note: For any Active/Active or Active/Passive DAG configuration it is important to size the redundant servers and storage with enough resources to be able to handle failover situations effectively.

5 Best Practice Recommendations

Each of the component layers in the system stack shown in Figure 22 requires careful design and configuration to ensure optimal system performance.

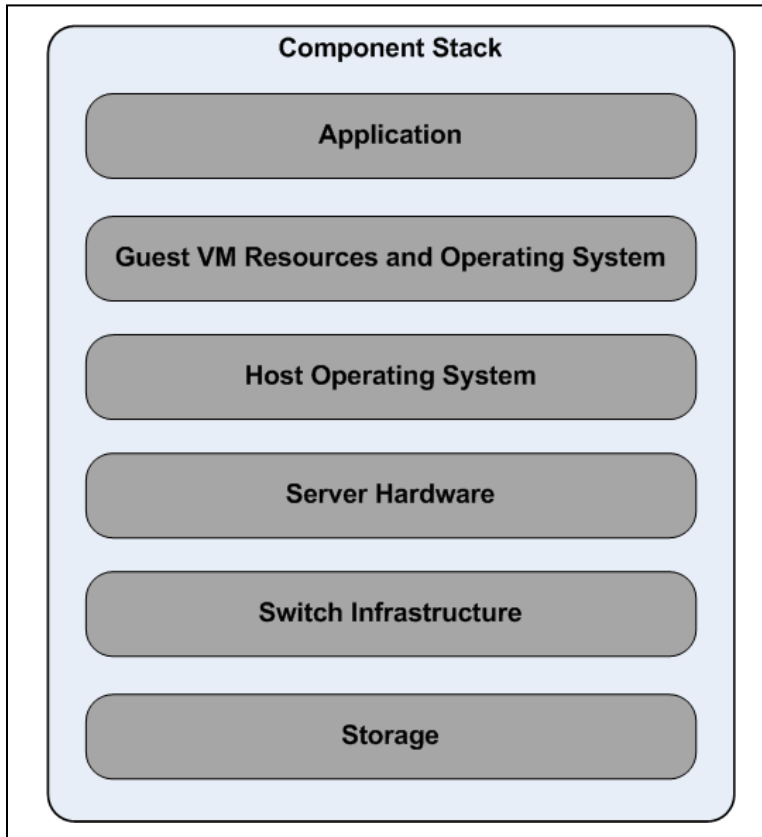


Figure 22 Component Stack

Within each layer of the stack you must consider each of the following design goals:

- Availability** Ensuring high availability is a critical design guideline for each component in the solution. Redundant sub-components and connection paths must be maintained to avoid single points of failure.
- Scalability** Scalability is also critical design goal. Your solution architecture should be able to scale to meet performance and capacity needs.
- Performance** Performance of each component in the stack must be balanced such that all performance requirements for the Exchange system are met (examples: I/O latency, user response times).

5.1 Storage

Considerations and recommendations for choosing the disk drives and RAID types are discussed in detail in Section 3. Additional recommendations for configuration of EqualLogic based storage for Exchange 2010 are as follows:

- It is a common best practice to implement Exchange based replication solution via a DAG to establish high levels of protection and recovery. If databases are deployed in a DAG, then hosting the database and log volumes on the same physical drives would be acceptable without compromising any recoverability requirements. However, in a DAG scenario we recommend as a best practice that you host active and passive database copies on separate disk drives. To support this using EqualLogic storage, you can store the active database and log volumes in one EqualLogic storage pool, with the passive or copy database and log volumes stored in a separate storage pool. The flexibility provided by the EqualLogic Peer Storage Architecture allows you to place active and passive copies of *different* databases in the same EqualLogic storage pool without compromising recoverability requirements. You should not place active and passive copies of the same database in the same pool.
- You should deploy separate volumes for databases and logs. Doing so will support isolation of respective entities for backup and storage based recovery options such as snapshots and clones.
- We recommend you host the EqualLogic storage volumes as Windows basic disks. Using GPT (GUID Partition Table) is recommended as the best practice for configuring partitions for Exchange Server 2010 volumes. Also it is a best practice to use NTFS allocation unit size of 64 KB.

Note: General recommendations for Exchange Mailbox storage design are provided by Microsoft® in the following document (Microsoft® TechNet Exchange Server TechCenter):

➤ *Mailbox Server Storage Design:*

<http://technet.microsoft.com/en-us/library/dd346703.aspx>

5.2 SAN Design

We recommend the following SAN design best practices:

- You should design redundant SAN component architectures. This includes the NICs on the servers and switches for the storage network (including server blade chassis switches and external switches).
- You should also make sure that the server NIC ports and storage NIC ports are connected in a way such that any single component failure in the SAN will not disable access to any storage array volumes.
- Flow control should be enabled on both the server NICs and the switch ports connecting to server and storage ports.
- Exchange Server 2010 storage I/O is primarily random in nature and uses a 32KB block size. However nightly database maintenance and backup operations create sequential I/O patterns. To best support this, we recommended you enable jumbo frames on the server ports and the switch ports. Within ESX you also need to set the jumbo frame MTU value on the virtual switches and the virtual NICs, including for the NIC device used by the guest VM OS.
- You should disable spanning tree on switch ports connecting to end devices like server ports and storage ports. Portfast should be enabled for these ports.

Note: General recommendations for EqualLogic PS Series array network configuration and performance is provided in the following document.

➤ *PS Series Array Network Performance Guidelines:*

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229>

5.3 Configuring Host Server Resources

It is important that you properly configure your servers to provide the optimal number of NICs for your Exchange I/O workload. You could consider the following:

- We recommend that you configure your servers to provide dedicated NICs for Exchange client/server I/O and separate dedicated NIC/HBAs for iSCSI SAN I/O.
- To provide for redundancy, you should provision at least two NICs per server for each path.
- If possible, Exchange mailbox server replication traffic should be segregated onto separate dedicated NICs. This will ensure that there is adequate connection bandwidth available for replication operations in a DAG and will increase replication performance where database changes caused by user operations are high.
- Additional NICs/HBAs should be provisioned for iSCSI SAN I/O as needed, based on throughput and IOPS requirements.

There are many resources for additional information that can help you configure your server resources:

Guidelines for Exchange 2010 hardware requirements in virtualized environments can be found here:

➤ *Exchange 2010 System Requirements*
<http://technet.microsoft.com/en-us/library/aa996719.aspx>

Guidelines for sizing Exchange server memory can be found here:

➤ *Understanding the Mailbox Database Cache*
<http://technet.microsoft.com/en-us/library/ee832793.aspx>

➤ *Understanding Memory Configurations and Exchange Performance*
<http://technet.microsoft.com/en-us/library/dd346700.aspx>

Guidelines for sizing Exchange server processor resources can be found here:

➤ *Mailbox Server Processor Capacity Planning*
<http://technet.microsoft.com/en-us/library/ee712771.aspx>

➤ *Understanding Processor Configurations and Exchange Performance*
<http://technet.microsoft.com/en-us/library/dd346699.aspx>.

Guidelines for resource ratios across different Exchange Server roles can be found here:

- *Understanding Server Role Ratios and Exchange Performance*
<http://technet.microsoft.com/en-us/library/dd346701.aspx>

Guidelines for resource ratios when combining multiple roles:

- *Understanding Multiple Server Role Configurations in Capacity Planning*
<http://technet.microsoft.com/en-us/library/dd298121.aspx>

Guidelines for sizing combined Hub and CAS roles in non-virtualized and virtualized environments can be found here:

- *Understanding Client Access and Hub Transport Combined Role Configurations in Capacity Planning*
<http://technet.microsoft.com/en-us/library/ee832795.aspx>

Additionally, there are many resources provided by Dell, including an Exchange 2010 Advisor tool, at the following location:

http://www.dell.com/content/topics/topic.aspx/global/shared/sitelets/solutions/messaging/en/exchange_2010

5.4 ESX Host Configuration

- If you choose to use the iSCSI software initiator in the ESX host, you should take advantage of EqualLogic aware connection and path management by installing the EqualLogic Multipathing Extension Module (MEM) for vSphere 4.1. See Section 5.5.
- Per Microsoft® Exchange Server storage requirements, Exchange Server guest OS's must be deployed on physical disk drives separate from physical drives hosting Exchange Server data.
- Host based availability solutions such as VMware HA, Fault Tolerance, vMotion, Distributed Resource Scheduling (DRS), etc. are not supported for the Mailbox Server role. Hence HA, DRS, and FT features should be turned off for any Exchange Mailbox Server guest virtual machines within ESX.
- You should configure separate virtual switches for VM network traffic and iSCSI storage traffic on the ESX hosts. If using both guest and host software initiators, it is a best practice to configure separate virtual switches for each connection path.
- Jumbo frames should be enabled on the iSCSI vSwitch. At least two server NICs dedicated for iSCSI traffic need to be configured as uplink NICs to the iSCSI vSwitch. If using the ESX software iSCSI initiator or iSCSI HBAs on the host, we recommended you install the EqualLogic MEM plug-in for VMware vSphere.

Hardware requirement guidelines for running Exchange 2010 in virtualized environments can be found here:

- *Exchange 2010 System Requirements*,
[http://technet.microsoft.com/en-us/library/aa996719\(printer\).aspx](http://technet.microsoft.com/en-us/library/aa996719(printer).aspx)

- Details on using the EqualLogic MEM plug-in can be found here:

<https://www.equallogic.com/support/download.aspx?id=1484>

5.5 EqualLogic Multipathing Extension Module for VMware vSphere 4.1

VMware vSphere offers many new and advanced enhancements to the software iSCSI initiator. The most significant of these enhancements is API support for third party multipathing plugins. The API provides a framework that enables the EqualLogic Multipathing Extension Module (MEM) to utilize VMware's Round Robin multipath management. Using the EqualLogic MEM, you can more intelligently route and efficiently load balance iSCSI traffic across multiple NICs. The EqualLogic MEM offers:

- Easy installation and configuration
- Automatic connection management
- Automatic load balancing across multiple active paths
- Increased bandwidth
- Reduced network latency
- Automatic failure detection and failover
- Multiple connections to a single iSCSI host

Note: Detailed steps for configuring MEM and the ESX host initiator with EqualLogic storage are provided in the following document:

- *Configuring and Installing the EqualLogic Multipathing Extension Module for VMware vSphere 4.1 and PS Series SANs:*

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=9823>

The iSCSI initiator you choose to use (the initiator in the ESX host or the initiator within the guest OS) will depend on factors such as backup and recovery process requirements and compatibility with other management tools used. If you use the iSCSI initiator within the guest OS (running Microsoft Windows 2008 Server), then you will be able to use EqualLogic Auto Snapshot Manager (ASM) / Microsoft Edition (ME). ASM/ME provides the ability to create consistent Exchange aware snapshots or clones of data volumes using storage hardware based snapshots or clones. It also enables other third party backup solutions based on Windows VSS (Volume Shadow copy Services) to utilize storage hardware based snapshots or clones for Exchange Server data backup. If you use the iSCSI initiator provided by the ESX host operating system, then other solutions should be investigated to provide for application consistent backups.

Even if you use the iSCSI initiator in the Windows guest VM for Exchange Server deployment, we recommend that you still use the ESX iSCSI initiator on the host, along with the EqualLogic MEM, for other virtual machines and applications that need it. In this case, we recommended you configure separate virtual switches for each storage type access path as follows:

- One vSwitch with its own set of NIC uplinks for iSCSI guest initiators
- One vSwitch with a second set of uplinks for the ESX software or hardware initiator

5.6 Virtual Machine and Guest OS Configuration

- You should allocate sufficient processor, memory and disk space resource to the guest VM and OS. Guidelines for processor and memory recommendations can be found here:
 - *Understanding Client Access and Hub Transport Combined Role Configurations in Capacity Planning*
<http://technet.microsoft.com/en-us/library/ee832795.aspx>.
- To take advantage of Exchange 2010 aware EqualLogic protection features such as Exchange consistent snapshots, clones and tools such as Auto-Snapshot Manager (ASM)/Microsoft® Edition, use the iSCSI initiator within the VM to connect to Exchange data and log volumes. Enable host software initiator along with EqualLogic MEM for other applications or virtual machines if using guest initiator.
- If using the iSCSI initiator within the guest OS, the following recommendations apply:
 - Create virtual NICs of type vmxnet3 within the guest VM to connect to iSCSI virtual switch.
 - Enable TSO and LRO in the guest VM NICs for iSCSI traffic.
 - We recommend you use the EqualLogic MPIO DSM installed as part of the EqualLogic Host Integration Toolkit (HIT Kit) in the guest OS.
- If you are using TCP Offload Engine (TOE) capable NICs, see the following VMware KB document for more details:
 - *What is a TOE NIC and does ESX support any TOE NICs?*
http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1006143
- More details on VMware support for TCP segmentation offload and jumbo frames can be found in the following vSphere Online Library document:
 - *TCP Segmentation Offload and Jumbo Frames*
http://pubs.vmware.com/vsp40/wwhelp/wwhimpl/js/html/wwhelp.htm#href=server_config/c_tcp_segmentation_offload_and_jumbo_frames.html
- A given VM cannot use more than one physical uplink on switch if it has only one virtual NIC. You should create as many virtual NICs as the number of uplinks to the iSCSI virtual switch. For more information on this see the following VMware Information Guide:
 - *VMware Virtual Network Concepts*
www.vmware.com/files/pdf/virtual_networking_concepts.pdf
- Install the latest version of VMware tools in the guest OS.
- Guidelines for virtualizing active directory domain controllers can be found in the following VMware publication:
 - *Virtualizing a Windows Active Directory Domain Infrastructure*
http://www.vmware.com/files/pdf/Virtualizing_Windows_Active_Directory.pdf

- It is a requirement to have Exchange Server guest OS and data hosted on separate physical drives. The guest OS system drive should be configured with at least 15GB plus the size of memory assigned to the guest VM to reserve disk space for any guest memory swap.
- We recommend you partition the volumes used for hosting Exchange server databases and logs as GPT disks. The volume should be formatted with NTFS with an allocation unit size of 64KB²⁰.

5.7 Exchange Server Configuration

- You should follow the guidelines provided in Section 0 for sizing Exchange Server mailbox database and log volumes.
- As a general best practice, we recommend you deploy database copies via a DAG for high availability and resiliency. In a DAG scenario, database and log volumes on a server can be hosted on the same RAID set (physical drives). However the respective copies need to be hosted on separate RAID sets or drives.
- We also recommend implementing a backup strategy that includes scheduled copying of databases to secondary storage. EqualLogic Auto Snapshot Manager (ASM/ME) can support Exchange recovery options using application consistent snapshots and clones. ASM/ME integrates with the Windows VSS framework to provide application consistent data copies on the EqualLogic SAN when used with the guest iSCSI initiator.

Guidelines for storage volume (LUN) layout to support Exchange 2010 can be found here:

- *Understanding Exchange 2010 LUN Architecture*
<http://technet.microsoft.com/en-us/library/ee832794.aspx>

²⁰ Reference: <http://technet.microsoft.com/en-us/library/ee832792.aspx#Best>

Appendix A Test System Components

The table below provides details for the major hardware and software system components used in the test system configuration.

Test Configuration - Hardware Components:	
Servers	<ul style="list-style-type: none"> • Dell PowerEdge M1000e Blade Enclosure • Dell PowerEdge M710 Blade Server: 2x Quad Core Intel® Xeon® X5570 Processors, 64 GB RAM, 8M Cache, 2.93 GHz, 2x Broadcom NetXtreme II 5709S Dual Port GbE Mezzanine Card installed
Network	<ul style="list-style-type: none"> • 2x Dell PowerConnect M6220 Blade Switches installed in M1000e blade enclosure fabrics A1 and A2 • 2x Dell PowerConnect M6348 Blade Switches installed in M1000e blade enclosure fabrics B1 and B2 • 2x Dell PowerConnect 6248 Switches installed in top of rack
Storage:	<ul style="list-style-type: none"> • Dell EqualLogic PS6000XV: 16 x 450GB 15K SAS disks with dual 4-port 1GbE controllers • Dell EqualLogic PS6000X: 16 x 600GB 10K SAS disks with dual 4-port 1GbE controllers • Dell EqualLogic PS6000E: 16 x 1024GB 7.2K SATA II disks with dual 4-port 1GbE controllers • Dell EqualLogic PS6500E: 46 x 1024GB 7.2K SATA II disks with dual 4-port 1GbE controllers
Test Configuration - Software Components:	
Operating systems:	<ul style="list-style-type: none"> • Host: VMware vSphere Server v4.1 Enterprise Plus (bare metal hypervisor directly installed on M710 blade server) <ul style="list-style-type: none"> ◦ MPIO enabled using EqualLogic Multi-Path Extension Module v1.0.0 for VMware vSphere 4.1 • Guest: Microsoft® Windows Server 2008 R2 Enterprise Edition (virtual machine): <ul style="list-style-type: none"> ◦ MPIO enabled using EqualLogic DSM for Windows when using guest iSCSI initiator ◦ EqualLogic Host Integration Toolkit v3.2.2 installed
Applications:	Microsoft® Exchange 2010 Enterprise Edition
Monitoring Tools :	<ul style="list-style-type: none"> • Windows Perfmon • EqualLogic SAN Headquarters
Simulation Tools :	<ul style="list-style-type: none"> • Microsoft® Exchange Server Jetstress 2010 Beta, version 14.00.0639.013 • Microsoft® Exchange Server Loadgen 2010 version 14.01.0180.003

Table 19 Test Configuration Components



THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN
TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS
IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.