

# **Microsoft Exchange 2010 Efficiency, Flexibility, Performance, and Availability at Scale Enabled by EMC Symmetrix VMAX, Virtual Provisioning, and VMware vSphere**

*A Detailed Review*

## **EMC Global Solutions**

---

### **Abstract**

This white paper demonstrates the performance of a large-scale, virtualized Microsoft Exchange Server 2010 configuration on EMC® Symmetrix® VMAX™ with Symmetrix Virtual Provisioning™ using VMware vSphere™. The document provides best practices and guidelines for simplifying an Exchange 2010 storage configuration using a building-block approach, and compares how both thick (regular, fully provisioned) and thin devices impact the Exchange 2010 environment.

May 2010

---

---

Copyright © 2010 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part number: H7018

---

## Table of Contents

---

Executive summary.....	5
Overview .....	5
Business case .....	6
Product solution.....	6
Key results.....	7
Introduction.....	8
Purpose .....	8
Scope .....	8
Audience .....	8
Terminology.....	9
Technology.....	10
Overview .....	10
EMC Symmetrix VMAX .....	10
EMC Symmetrix Virtual Provisioning .....	10
VMware vSphere.....	11
Environment .....	12
Simulated Exchange 2010 configuration .....	12
Hardware.....	13
Software .....	13
Physical environment .....	14
Exchange Server 2010 storage design guidance on the Symmetrix VMAX.....	15
Overview .....	15
Design guidelines for virtualized Exchange 2010 on the Symmetrix VMAX array .....	15
Sizing storage for the Exchange 2010 environment .....	16
What is a building-block? .....	16
Exchange 2010 building-block design process .....	17
Applying the building-block design process to Exchange 2010.....	18
Storage requirements calculated for the Exchange 2010 test environment .....	19
Building-block design .....	20
Building-block for the virtualized environment .....	21
Symmetrix Virtual Provisioning with Exchange 2010.....	22
Overview .....	22
Design guidelines for using Symmetrix Virtual Provisioning with Exchange 2010 .....	22
Testing methodology and results .....	23
Overview .....	23

---

Disk layout validation testing with Jetstress .....	23
Introduction.....	23
Microsoft Jetstress 2010 .....	23
Jetstress testing approach .....	23
Jetstress test configuration .....	24
Jetstress test types.....	24
Jetstress storage performance results .....	25
SAN performance during testing .....	27
Symmetrix VMAX I/O throughput .....	28
Database backup and recovery performance overview .....	29
Backup VSS test: Database Read-only performance .....	29
Soft Recovery test: Transaction log recovery/replay performance .....	29
Validating the performance impact of thin versus regular devices with Jetstress .....	30
Introduction.....	30
Virtual Provisioning testing approach.....	30
Test results—Comparing thin devices versus regular devices .....	31
Test results—Comparing thin device latencies versus regular device latencies .....	32
vSphere mailbox server validation testing with LoadGen .....	33
Introduction.....	33
Microsoft LoadGen .....	33
LoadGen testing approach.....	33
LoadGen test configuration .....	33
LoadGen mail client performance results .....	33
LoadGen VM CPU utilization results.....	34
Conclusion.....	35
Summary .....	35
Findings .....	35
References .....	36
White papers .....	36
Product documentation .....	36
Other documentation.....	36

---

---

## Executive summary

---

### Overview

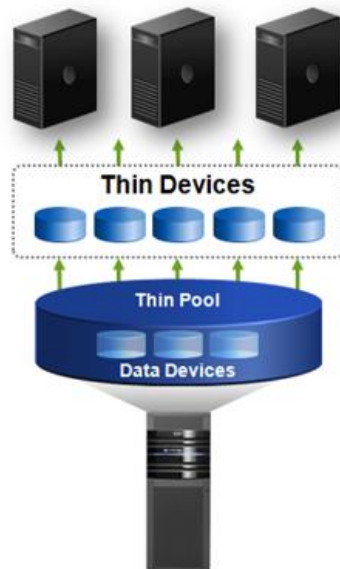
Large, multifaceted organizations need to deploy a 100 percent reliable messaging infrastructure that can lower communication costs, increase productivity, simplify administration, and decrease IT overhead. The introduction of Microsoft Exchange Server 2010 allows enterprises to tailor their existing Exchange infrastructure to specific organizational business and technology needs, while lowering overall storage costs.

In anticipation of the application's full support of much larger user mailboxes (10 GB) IT professionals are looking for a way to manage the potential shift in storage requirements associated with a surge in Exchange data. As user mailboxes experience rapid growth Exchange Server 2010 performance must be balanced with storage capacity to achieve acceptable cost and performance metrics.

This white paper addresses how new levels of manageability and reliability can be achieved by deploying Exchange Server 2010 with the full-featured EMC<sup>®</sup> Symmetrix<sup>®</sup> VMAX<sup>™</sup> array integrating powerful Symmetrix VMAX Engines for storage and consolidation. The document also demonstrates how using VMware vSphere<sup>™</sup> 4 virtualization technology can greatly reduce your Exchange 2010 footprint, energy consumption, and overall cost in administering Exchange 2010.

Within this white paper, we will provide information gained through extensive testing of thin provisioning provided by the EMC Symmetrix Virtual Provisioning<sup>™</sup> feature. Virtual Provisioning allows the Symmetrix array to be provisioned in large pools of storage so that the space can be allocated thick (fully formatting the entire drive) or thin (only consuming space as needed).

As illustrated below, thin devices as provided by Virtual Provisioning enable organizations to present a virtual amount of capacity to a host, allowing the application to consume space only as needed without reallocation or reconfiguration. This lowers total cost of ownership (TCO) by reducing initial allocation of storage capacity and simplifies management.



---

This white paper discusses:

- A step-by-step design methodology for configuring virtualized Exchange 2010 on the Symmetrix VMAX using a basic building-block approach based on a very heavy user profile
- Symmetrix VMAX storage performance validation for Exchange 2010 IOPS
- How to achieve increased cost savings in the Exchange 2010 environment through server hardware consolidation enabled by VMware vSphere
- How Symmetrix thin devices (part of Virtual Provisioning) provide clear resource savings, in addition to efficiencies in allocating space and provisioning storage

---

**Business case** Enterprise organizations need a simple plan for integrating Exchange Server 2010 into their existing data centers. This could be characterized as a complicated task without EMC guidance and real-world, validated test results.

By adopting this Exchange 2010 design approach, customers will achieve key business objectives and learn how to:

- Use a simple building-block approach for server and storage design for Exchange 2010
- Minimize the initial cost of storage deployment using Symmetrix Virtual Provisioning (customers do not have to purchase the entire storage requirements for future storage during the deployment stage)
- Reduce operational and capital expenses that result from information growth and unused resources with VMware vSphere virtualization software
- Reduce long-term storage costs
- Improve e-mail server performance and IT productivity

---

**Product solution** This white paper presents information on the reliable performance of the Symmetrix VMAX array with Microsoft Exchange 2010 by consolidating a large amount of users in a single, virtualized environment.

The test configuration is designed to support two database copies for 120,000 very heavy (very heavy as defined by Microsoft as .18 IOPS) Exchange 2010 users on a single Symmetrix VMAX array. The simulated test environment implements a very reliable SAN infrastructure and VMware vSphere virtual mailbox configuration. The mailbox configuration consists of 24 virtualized Exchange 2010 Mailbox Servers housed on just six physical servers. By adopting this design, customers can realize a 4:1 server savings ratio.

With Exchange Server 2010, Microsoft is promoting the idea of much larger user mailboxes (10 GB). This may cause additional challenges to arise such as a large increase in storage requirements, which may require customers to rework the existing backup and restore process. To address storage challenges, the powerful Symmetrix VMAX array is used to provide all the necessary performance, maximum data capacity, and multiple level data protection for Exchange Server 2010.

Symmetrix Virtual Provisioning is used to alleviate one of the biggest problems facing Exchange storage administrators: Overprovisioning. Overprovisioning of physical storage translates to an increase in power consumption, cooling and floor

---

space requirements. Thin provisioning solves this problem by allowing more storage to be presented to an application than is physically available.

More importantly, thin provisioning allocates physical storage only when the storage is actually written to. This allows administrators more flexibility in estimating future growth, and reduces the initial cost of provisioning storage dedicated to Exchange 2010.

An additional cost saving measure in our design was to use VMware vSphere, which reduces IT costs, improves flexibility and provides for physical server consolidation. VMware vSphere also increases energy efficiency by running fewer servers.

---

## **Key results**

The following list details the key results gathered during Exchange Server 2010 solution testing:

- The virtualized Exchange 2010 building-block approach used to design the test environment is easy to use and scales well, producing excellent performance results.
- A 4:1 server consolidation ratio can be achieved by incorporating VMware vSphere as the server virtualization platform.
- The Symmetrix VMAX array is capable of producing many more Exchange 2010 IOPS per spindle in comparison to DAS configurations. This will result in fewer spindles required to produce the performance required.
- Thin devices, part of the Symmetrix Virtual Provisioning technology, allow customers to purchase only the storage required for the initial mailbox size. As the user mailboxes grow, more storage can seamlessly be added with no effect to the users or Exchange server's performance.

For more information, refer to the [Testing methodology and results](#) section.

---

---

## Introduction

---

**Purpose** The purpose of this white paper is to provide performance results, storage design guidelines, and best practices for designing virtualized Symmetrix VMAX storage with Exchange Server 2010.

---

**Scope** The scope of this white paper is to document:

- Design guidelines used to build this Exchange 2010 environment
- Exchange 2010 performance test results
- Symmetrix VMAX storage configuration test results using the following tools:
  - Microsoft Jetstress 2010
  - Microsoft Load Generator 2010 (LoadGen)
- How to use VMware vSphere to virtualize Exchange 2010 mailbox servers
- Benefits of Symmetrix Virtual Provisioning with Exchange 2010 data

---

**Audience** This white paper is intended for:

- EMC employees
- EMC partners
- Customers, including IT planners, storage architects, and administrators
- Field personnel, who are tasked with designing and implementing Microsoft Exchange Server 2010 solutions with the Symmetrix VMAX storage array

---

---

## Terminology

This white paper includes the following terms.

Term	Definition
<b>Background Database Maintenance (BDM)</b>	Process of the Exchange database maintenance that involves checksumming both active and passive database copies.
<b>Database Availability Group (DAG)</b>	The new Exchange 2010 base component that provides high availability (HA) and site resilience. A DAG can contain up to 16 mailbox servers. The servers host a set of databases that provide automatic database-level recovery from failures affecting individual databases. Any server in a DAG can host a copy of a mailbox database from any other server in the DAG.
<b>Regular devices</b>	Represent fully-provisioned, traditional LUNs (also known as “thick” devices). In this conventional storage provisioning model, space is allocated in excess of current storage requirements in anticipation of expected growth.
<b>Thin devices</b>	A host accessible device that has no storage directly associated with it.
<b>Thin LUNs</b>	A logical unit of storage where physical space allocated on the storage system may be less than the user capacity seen by the host server.
<b>Thin pool</b>	A group of disk drives used specifically by thin LUNs. There may be zero or more thin pools on a system. Disks may be a member of no more than one thin pool.
<b>Thin provisioning</b>	<p>An industry standard technology whereby a large “thin device” (that is, a volume) is configured and presented to the host while consuming physical storage from a shared pool only as needed.</p> <p>Thin provisioning, provided by EMC's Virtual Provisioning technology, relies on the on-demand allocation of blocks of data versus the traditional method of allocating all the blocks up front. This methodology eliminates almost all whitespace, which helps avoid the poor utilization rates that occur in the traditional storage allocation method (thick provisioning) where large pools of storage capacity are allocated to individual servers but remain unused (not written to).</p>
<b>Virtual Provisioning</b>	Virtual Provisioning is EMC's implementation of thin provisioning, which enables organizations to present a certain amount of virtual capacity to a host, allowing the application to consume space as needed, for both thin and thick (fully provisioned) pool LUNs. This technology lowers total cost of ownership (TCO) and simplifies management by reducing administrative tasks required to support growth.

---

## Technology

---

### Overview

The following sections identify and briefly describe the technology and components used in the configuration, including:

- EMC Symmetrix VMAX
  - EMC Symmetrix Virtual Provisioning
  - Microsoft Exchange Server 2010
  - VMware vSphere
- 

### EMC Symmetrix VMAX

The Symmetrix VMAX is built on the strategy of simple, intelligent, modular storage and incorporates a new Virtual Matrix™ interface that connects and shares resources across all nodes. This allows the storage array to seamlessly grow from an entry-level configuration into the world's largest storage system. The storage array provides the highest levels of performance and availability featuring:

- Up to 2 petabytes (PB) usable capacity
  - Up to 128 Fibre Channel (FC) front-end (FE) ports for host connections
  - FE ports
  - Up to 64 FICON FE ports
  - Up to 64 GbE / iSCSI FE ports
  - Up to 1 terabyte (TB) global memory (512 GB usable)
  - 48 to 2,400 drives
  - Enterprise Flash Drives (EFDs), 200/400 GB
  - FC drives
    - 146/300/450 GB, 15k rpm
    - 300/400/450/600 GB, 10k rpm
  - SATA drives, 1 TB, 7.2k rpm
- 

### EMC Symmetrix Virtual Provisioning

An integrated Symmetrix array feature that enables organizations to present a certain amount of virtual capacity to a host, allowing the application to consume space as needed, for both thin and thick (fully provisioned) pool LUNs. This technology lowers TCO by reducing initial allocation of storage capacity and simplifies management by reducing administrative tasks required to support growth.

---

---

**Microsoft Exchange Server 2010 Enterprise Edition**

Microsoft Exchange Server 2010 is an enterprise e-mail and communication system that allows businesses and customers to collaborate and share information. EMC enhances Exchange Server 2010 with the industry's broadest choice of storage platforms, software, and services.

With the new version of Exchange 2010 Microsoft presents a number of new features including, but not limited to:

- Database Availability Groups for database high availability
- Online mailbox moves
- Larger mailboxes (10 GB mailboxes)

Mailbox servers can now be implemented in database resiliency configurations with database-level replication and failover. Major improvements with the application database structure and I/O reduction include support for a larger variety of disk and RAID configurations including FC and SATA drives.

---

**VMware vSphere**

VMware vSphere is the industry's first cloud operating system, transforming data centers into a dramatically simplified cloud infrastructure and enabling the next generation of flexible, reliable IT services. Leveraging VMware's industry-leading technology and experience, VMware vSphere delivers uncompromised control, in the most efficient manner, while fully preserving customer choice.

---

---

## Environment

---

### Simulated Exchange 2010 configuration

The following table provides additional details about the simulated Exchange 2010 test environment.

Item	Value
Number of Exchange mailboxes simulated	120,000 mailboxes
Number of ESX servers	6 (physical servers)
Number of users/virtual machines (VMs) per ESX server	20,000 users and four VMs
Total number of Exchange VMs	24 Exchange VM mailbox servers with 5,000 users per VM
Number of DAGs	3
Number of servers/DAG	16
Number of active mailboxes/server	5,000
Number of databases/hosts	8
Number of database copies (Mailbox Resiliency)	2
Number of mailboxes/database	625
Simulated profile: I/Os per second per mailbox (IOPS, includes 20% headroom)	.18
Mailbox size	1 GB
Database Read/Write ratio (in Mailbox Resiliency configuration)	3:2
Database maintenance configuration	24x7 BDM
Database LUN size	900 GB
Log LUN size	80 GB
Total database size for performance testing	618 GB (DB) 99.4 GB (LUN)
Percent storage capacity used by Exchange database	77%

---

---

**Hardware**

The hardware used in the environment is listed below.

Item	Description
EMC Symmetrix VMAX	Microcode version 5874.207
Storage connectivity (FC, SAS, SATA, iSCSI)	FC
Storage cache	240, 640 MB
Number of storage controllers	4
Number of storage ports used	12
Maximum bandwidth of storage connectivity to host	12 * 4 Gb/s
FC switch	4 Gb/s enterprise class FC switch
Host Bus Adapter (HBA) and firmware	2 Dual port, 4 GB HBAs (QLA2562)
Number of HBAs/host	2
Host servers	Intel Xeon CPU X7350 @ 2.93 GHz, 4 Core, 128 MB memory
Total number of disks tested in the solution	384, 600 GB, 10k rpm drives
Maximum number of spindles can be hosted in the storage	2,400

---

**Software**

The software used in the environment is listed below.

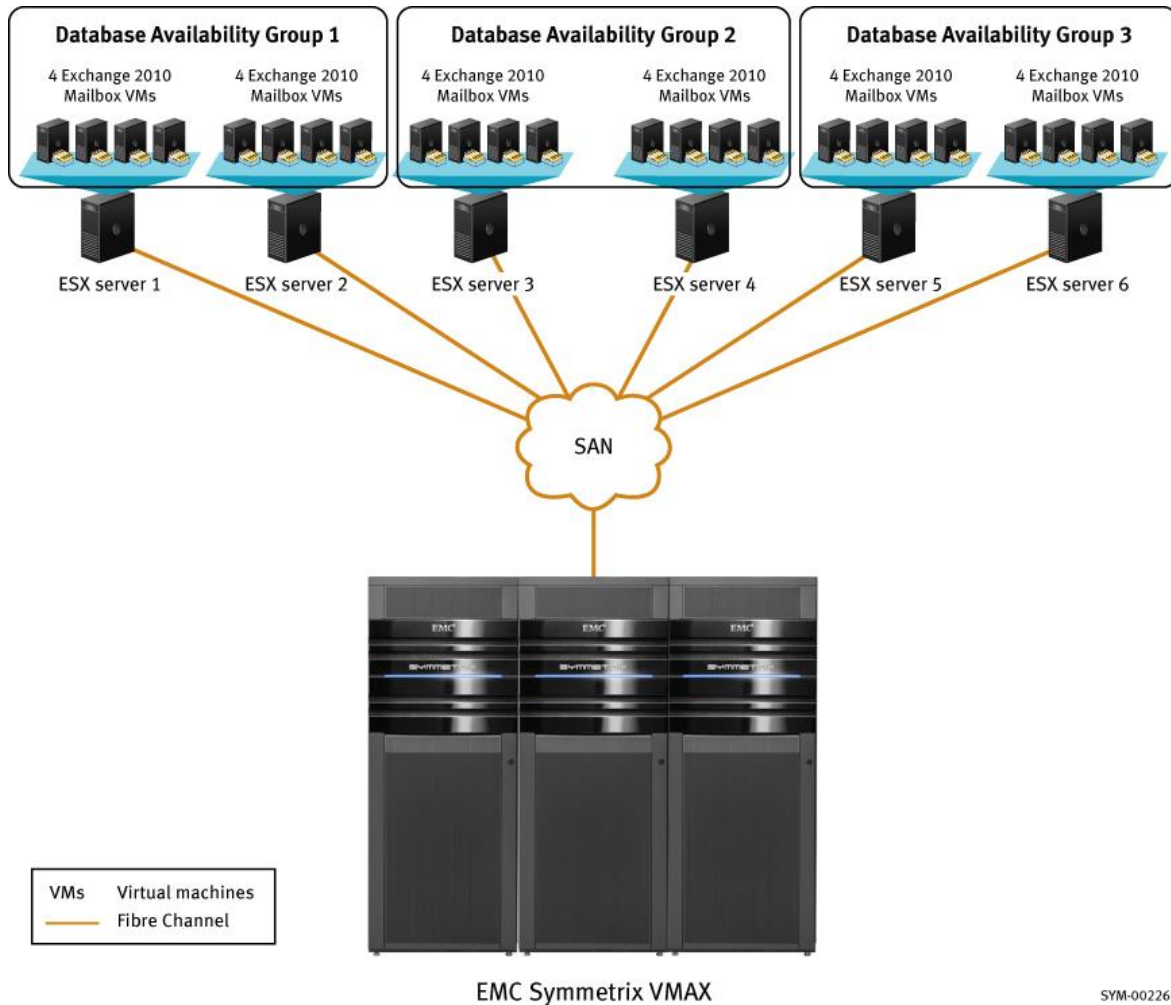
Item	Description
HBA driver	QLogic 9.1.7.16 2/15/2008
HBA QueueTarget setting	256
HBA QueueDepth setting	256
Multipathing	EMC PowerPath®/VE 5.4 (64 bit)
Host OS	Microsoft Windows Server 2008 Enterprise Service Pack 2
Microsoft Exchange Server 2010 Enterprise Edition	14.0.639.19
VMware vSphere	Version 4.0, Update 1, Build 208111

**Physical environment**

The following diagram depicts the physical architecture of the simulated test environment.

**Note**

In production, the environment would need to be doubled to accommodate the full DAG setup. Our testing represents a worst case scenario where half the environment is down and all databases are active on each server.



---

## Exchange Server 2010 storage design guidance on the Symmetrix VMAX

---

### Overview

The following sections provide design guidance for creating a virtualized Exchange 2010 building-block on the Symmetrix VMAX. The process is essentially the same as designing a physical environment, except that a physical server is required for each VM.

---

### Design guidelines for virtualized Exchange 2010 on the Symmetrix VMAX array

The following list details the storage sizing guidelines followed during solution testing.

- Always calculate I/O spindle requirements first, then capacity requirements.
  - I/O requirements include user I/O (send\receive) plus 20 percent. Sequential I/Os like log and BDM are not factored in. See *Microsoft TechNet* guidance below:

Pure sequential I/O operations are not factored into the IOPS per Mailbox calculation since storage subsystems can handle sequential I/O much more efficiently than random I/O. These operations include BDM, log transactional I/O, and log replication I/O.
  - Database and log I/O should be evenly distributed among the Symmetrix VMAX Engines.
  - Use the Exchange building-block approach whenever possible for both regular (thick) and thin devices.
  - To achieve better performance, use fewer, larger hypervolumes to create meta LUNs.
  - A minimum of two HBAs are required per server (one connected to an even director and the other connected to an odd director).
  - A minimum of three IP connections are preferred for each ESX server.
  - Use the default Read/Write ratio 3:2 with Exchange 2010.
  - Isolate the Microsoft Exchange server database workload from other I/O-intensive applications or workloads. This ensures the highest level of performance for Exchange and simplifies troubleshooting in the event of a disk-related Microsoft Exchange performance issue.
  - Install EMC PowerPath/VE for optimal path management and maximum I/O performance. For more information on installing and configuring the PowerPath application, visit:  
<http://www.emc.com/products/detail/software/powerpath-multipathing.htm>
  - Visit the following Microsoft websites for guidance on determining server memory and CPU requirements:
    - <http://technet.microsoft.com/en-us/library/ee832793.aspx> (Memory)
    - <http://technet.microsoft.com/en-us/library/ee712771.aspx> (CPU)
-

---

**Sizing storage for the Exchange 2010 environment**

Sizing and configuring storage for use with Microsoft Exchange Server 2010 can be a complicated process, driven by many variables and factors, which vary from organization to organization. Properly configured Exchange storage, combined with properly sized server and network infrastructure, can guarantee smooth Exchange operation and best user experience.

One of the methods that can be used to simplify the sizing and configuration of large Microsoft Exchange Server 2010 environments is to define a unit of measure—a *building-block*.

---

**What is a building-block?**

A building-block represents the required amount of disk and server resources required to support a specified number of Exchange 2010 users. The amount of required resources is derived from:

- A specific user profile type
- Mailbox size
- Disk requirements

Using the building-block approach takes out the guesswork and simplifies the implementation of Exchange Server 2010. Once the initial building-block is designed, it can be easily reproduced to support the required number of total users in your organization.

By using this approach, Exchange administrators can now create their own building-blocks that are based on their company's specific Exchange environment requirements. This approach is very helpful when future growth is expected, as it makes Exchange environment expansion much easier, and straightforward.

EMC's best practices involving the building-block approach for Exchange Server design proved to be very successful throughout many customer implementations.

---

**Exchange 2010 building-block design process**

The high-level building-block design process for Exchange 2010 is similar to that used for Exchange 2007. Review the following table for an understanding of the process flow used to develop and validate the test environment's storage design.

Phase	Description
1	<p><b><u>Collect user requirements</u></b></p> <p>In this phase, the Exchange administrator identifies:</p> <ul style="list-style-type: none"> <li>• Number of users</li> <li>• User I/O profile = Send\Receive and message size per user</li> <li>• Mailbox size</li> <li>• Deleted item retention</li> <li>• Concurrency</li> <li>• Replication required, number of DAG database copies</li> <li>• Backup/Restore requirements (RTO or RPO)</li> <li>• Third-party software that effects space or I/O (BlackBerry, antivirus software)</li> </ul>
2	<p><b><u>Design the storage architecture based on user requirements</u></b></p> <p>In this phase, the Exchange design is developed using the following tools:</p> <ul style="list-style-type: none"> <li>• EMC Exchange building-block methodology</li> <li>• EMC Exchange 2010 solutions</li> <li>• EMC and Microsoft best practices</li> <li>• Published Exchange 2010 Exchange Solution Review Program (ESRP) documentation:  <a href="http://technet.microsoft.com/en-us/exchange/ff182054.aspx">http://technet.microsoft.com/en-us/exchange/ff182054.aspx</a>.</li> </ul>
3	<p><b><u>Validate the design</u></b></p> <p>In this phase, the design is validated with the following tools:</p> <ul style="list-style-type: none"> <li>• Jetstress is used to validate storage</li> <li>• LoadGen is used to validate storage and Exchange server performance</li> </ul>

**Applying the building-block design process to Exchange 2010**

The following sections walk you through the storage design process applied.

**Phase 1—Collect user requirements**

The user requirements used to validate both the building-block storage design methodology and Symmetrix VMAX performance are detailed in the following table.

Item	User requirement
Number of users per server	5,000
User I/O profile	Very heavy 30 sent\120 received + 20% = .18 IOPS per user
Mailbox size	1 GB plus
Deleted item retention	14 days
Concurrency	100%
Replication required	Yes (2 DAG copies, 1 Active/1 Passive)
Backup/Restore required	Yes (VSS clone)

**Phase 2—Design the storage architecture based on user requirements**

The following basic formula can be used to calculate storage for the Exchange 2010 database and logs.

$$(IOPS * \%R) + WP (IOPS * \%W) / \text{Physical Disk Speed} = \text{Required Physical Disks}$$

<b>Where</b>	<b>Is</b>
<b>IOPS</b>	the number of input/output operations per second
<b>%R</b>	the percentage of I/Os that are reads
<b>%W</b>	the percentage of I/Os that are writes
<b>WP</b>	the RAID write penalty multiplier (RAID 1=2, RAID 5=4)
<b>Physical Disk Speed</b>	140 for 10k rpm drives, or 65 for 7.2k rpm SATA drives

**Important**

First calculate the user IOPS for each building-block. **(Users \* IOPS per user) + 20% (overhead and spikes)**. Next, add in the write penalty and divide by I/O per spindle for Exchange 2010.

**Phase 3—Validate the design**

For a complete summary of Jetstress and LoadGen findings, see the [Testing methodology and results](#) section.

---

**Storage requirements calculated for the Exchange 2010 test environment**

The calculations for the Exchange 2010 test environment are summarized below. These requirements were calculated using the formula detailed in [Exchange Server 2010 storage design guidance on the Symmetrix VMAX>Applying the building-block design process to Exchange 2010](#).

**IOPS Requirements**

- 5,000 users \* .18 = 900 + 20% = 1,080 IOPS
- RAID 5 FC: R (1,080 \*.60) + W 4(1,080 \*.40) = 2,376 / 140 = **16** disks
- RAID 1/0 FC: R (1,080 \*.60) + W 2(1,080 \*.40) = 1,512 / 140 = **10** disks
- RAID 1/0 SATA: R (1,080 \*.60) + W 2(1,080\*.40) = 1,512 / 65 = **22** disks

**Capacity Requirements**

- DB Space: 5,000 users \* 1 GB = 5,000 GB \* 35% = 6,750 GB
- Log space: Very heavy users with 75k avg message = 39 logs per day, based on MS calculation, 39 \* 4 days (log retention buffer) = 156 \* 5,000 = 780 GB
- Total space required: **7,530** GB for database and log LUNs

**Disk space based on RAID and disk type**

- Sixteen 600 GB 10k rpm drives in two RAID 5 (7+1) DG = **7,630** GB
- Ten 600 GB 10k rpm drives in RAID 1\0 = **2,880** GB, require **24** in RAID 1/0 configuration

**Number of Disks Required per DB copy**

Based on the above calculations, sixteen 600 GB 10k rpm disks are the best choice to support the 5,000-user building-block. With RAID 1/0 24 10k drives are required from a space perspective. With SATA 22 drives are required from an I/O perspective.

---

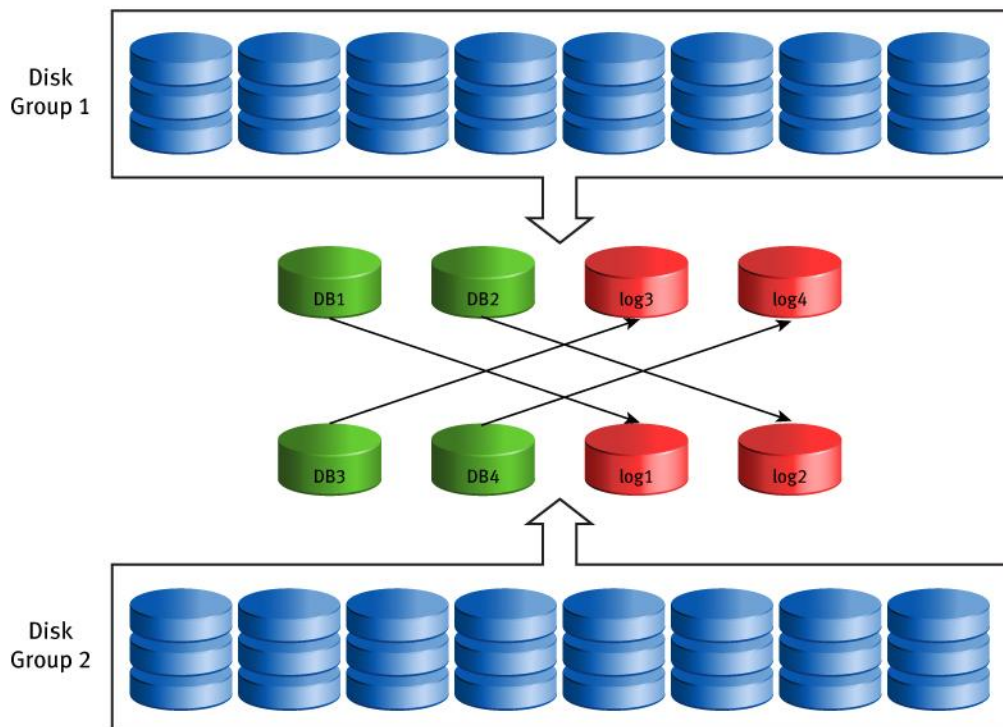
---

## Building-block design

The next step is to identify how many databases to configure per Exchange server and how the disk groups will be laid out. The following image shows one suggested design based on EMC and Microsoft recommended guidelines for Exchange 2010.

### VMAX Exchange 2010 building-block

- 600 GB 10k FC drives in 2 R5 (7+1)
- 7,630 GB available storage per building-block
- 5,000 Users with 1 GB Mailbox & .18 IOPS
- 16 LUNs (8 DB, 8 Log)
- DB LUN - 900 GB
- Log LUN - 80 GB
- 8 Way Metas



SYM-002501

As shown in the image, four Exchange databases and four log LUNs share each 7+1 disk group, but each Exchange database and its corresponding log LUNs are placed on separate disk groups. Database and log devices were presented to the VMs as RDM in physical compatibility mode and configured to support:

625 users, with 1 GB mailboxes

#### Note

See [Symmetrix VMAX Virtual Provisioning with Exchange 2010](#) in this white paper for information on how the Symmetrix VMAX Virtual Provisioning feature can easily increase the user mailbox limit with no impact to the users or servers.

---

**Building-block  
for the  
virtualized  
environment**

Determining the correct VM memory and CPU requirements are very important when creating an Exchange building-block. The CPU and memory requirements are based on Microsoft guidance from Microsoft TechNet:

<http://technet.microsoft.com/en-us/library/ee712771.aspx>

The following table summarizes the building-block for the VMware environment created for performance testing.

Item	Description
Number of users supported	5,000
User profile supported	.18 IOPS "very heavy" 150 messages/user/day (30 sent/120 received, plus 20%)
Mailbox size	1 GB simulated
RAID type	RAID 5 (7+1)
Database LUN size	900 GB (625 users per database)
Log LUN size	80 GB
VM CPU	6 CPU Xeon X7460 2.66 GHz
VM Memory	32 GB
Total disks required	Sixteen 600 GB FC disks, 10k rpm

---

## Symmetrix Virtual Provisioning with Exchange 2010

### Overview

Exchange 2010 allows for and encourages large user mailboxes (10 GB) and large databases (up to 2 TB). When moving from a 200 MB user mailbox in Exchange 2003 or 2007 to a 10 GB mailbox in 2010 it could take years for user mailboxes to reach 10 GB.

The Symmetrix Virtual Provisioning feature allows for the addition of storage space as a user's mailbox grows. It is possible to start with storage to accommodate 1 or 2 GB mailbox per user, and add storage incrementally depending on user need.

From an economic standpoint, Symmetrix Virtual Provisioning can save customer money on:

- Power and cooling
- Floor space
- Purchasing additional disks

Another advantage of Symmetrix Virtual Provisioning is it allows for more storage to be presented to the server than is physically available. For example, an Exchange mailbox server may see a 2 TB LUN for a particular database but only the storage space being used by the database will be accessed. With Virtual Provisioning, when storage is added and used by a database it is invisible to the host.

It is important to note that, like a regular LUN, virtual provisioned thin LUNs are compatible with the same technologies including VMware, Snaps, and Clones.

For a complete feature overview and best practice recommendations for deploying VMware vSphere on the EMC Symmetrix VMAX using Virtual Provisioning, see the following white paper available at <http://powerlink.emc.com>:

*Implementing EMC Symmetrix Virtual Provisioning with VMware vSphere—Applied Technology white paper*

---

### Design guidelines for using Symmetrix Virtual Provisioning with Exchange 2010

The following summary details Symmetrix Virtual Provisioning guidelines that should be considered when designing Exchange layouts on thin LUNs.

- The initial disk devices in the thin pool need to be able to handle the I/O performance requirements.
  - Dedicate a thin pool per Exchange 2010 mailbox server to provide for more granular administration and monitoring.
  - It is also possible to create larger thin pools that contain the storage for multiple Exchange mailbox servers.
  - Use large data devices for your thin pool.
  - Monitor the thin pool for available space using EMC Symmetrix Management Console (SMC) alerts, EMC Ionix™ ControlCenter® (ECC), or EMC Solutions Enabler.
-

---

## Testing methodology and results

---

### Overview

The following sections prove the efficiency of Exchange 2010 operations in the virtualized environment using three testing phases as follows:

- Validating the disk design methodology and the 5,000-user building-block
  - Validating the building-block's scalability across the Symmetrix VMAX and ESX servers
  - Determining the optimal Symmetrix VMAX configuration for Exchange 2010
  - Validating Symmetrix Virtual Provisioning with Exchange 2010
- 

## Disk layout validation testing with Jetstress

---

### Introduction

This section details disk layout analysis in the Exchange 2010 test environment.

---

### Microsoft Jetstress 2010

The Jetstress tool verifies the performance and stability of a disk subsystem prior to putting an Exchange 2010 server into production. Jetstress helps verify disk performance by simulating Exchange disk I/O load. Specifically, Jetstress simulates the Exchange database and log file loads produced by a specified number of users.

For more information on Jetstress, visit:

<http://msexchangeteam.com/archive/2009/09/01/452271.aspx>

---

### Jetstress testing approach

Jetstress 2010 was used to verify the performance and stability of the Symmetrix VMAX array. Jetstress helps verify disk performance by simulating Exchange disk I/O load. Specifically, Jetstress simulates the Exchange database and log file loads produced by a specified number of users. It is also important to note that Jetstress testing focuses on storage solution testing, and highlights performance and reliability issues with storage design.

#### Note

The Jetstress tool is designed to test performance of the Exchange storage subsystem before placing it in the production environment. It is not designed to test server CPU and memory configuration and impact of MAPI user activity. The Microsoft LoadGen tool is used to server CPU and memory configuration.

---

---

**Jetstress test configuration**

Jetstress test results are provided based on the following configuration:

- 120,000 Very Heavy users (with BDM running)
- 6 ESX servers

After successfully completing the Jetstress disk performance and stress tests in a non-production environment, you will have ensured that your Exchange 2010 disk subsystem is adequately sized (in terms of performance criteria you establish) for your specific user count and profiles.

---

**Jetstress test types**

Four types of Jetstress tests were run in the test environment to verify storage reliability as detailed in the following table.

Test	Test name	Description
1	Two-hour performance test	Measures baseline storage I/O performance and how long it takes for the storage to respond to an I/O under load.
2	24-hour stress test	Validates how storage responds to a high I/O load for an extended period of time.
3	Database backup VSS test	Measures the maximum rate at which databases can be backed up via VSS.
4	Soft recovery test	Measures the maximum rate at which the log files can be played against the databases.

---

**Jetstress  
storage  
performance  
results**

Primary storage performance testing is designed to exercise storage with the maximum Exchange type of I/O for a period of two hours. The test is designed to show how long it takes for the storage to respond to an I/O under load.

Jetstress results for each server were very similar and well balanced. The results are presented in the following tables, and categorized by:

- Mailbox VM
- ESX server, and
- Total results for the environment (with 120,000 users)

**Mailbox VM Performance Results (5,000 users)**

<b>Database I/Os</b>	<b>Avg performance per Mailbox VM (5,000 users)</b>
DB disk transfers/s	929
DB disk Reads/s	553
DB disk Writes/s	376
Avg DB disk Read latency	16.61 ms
Avg DB disk Write latency	9.72 ms
<b>Transactional log I/Os</b>	
Log disk Writes/s	309
Avg log disk Write latency	2.69
<b>BDM IOPS</b>	
Total BDM IOPS	233 IOPS
Total IOPS per Mailbox VM	1,471 IOPS

**ESX server Performance Results (20,000 users)**

<b>Database I/Os</b>	<b>Avg performance per ESX server (4 Mailbox VMs/20,000 users)</b>
DB disk transfers/s	3,716
DB disk Reads/s	2,212
DB disk Writes/s	1,504
Avg DB disk Read latency	16.61 ms
Avg DB disk Write latency	9.72 ms
<b>Transactional log I/Os</b>	
Log disk Writes/s	1,236
Avg log disk Write latency	2.69
<b>BDM IOPS</b>	
Total BDM IOPS	932 (256k Reads)
Total IOPS per Mailbox VM	5,884 IOPS

---

**Total performance results for six ESX servers (120,000 users)**

<b>Database I/Os</b>	<b>Total performance results for six ESX servers (24 Mailbox VMs/120,000 users)</b>
DB disk transfers/s	22,296
DB disk Reads/s	13,272
DB disk Writes/s	9,024
Avg DB disk Read latency	16.61 ms
Avg DB disk Write latency	9.72 ms
<b>Transactional log I/Os</b>	
Log disk Writes/s	7,416
Avg log disk Write latency	2.69
<b>BDM IOPS</b>	
Total BDM IOPS	5,592 (256k Reads)
<b>Total IOPS</b>	<b>35,304 IOPS</b>

---

**SAN performance during testing**

From a SAN and storage performance perspective it is important that the I/O is balanced across the HBAs, FC switch ports, and Symmetrix VMAX Fibre Adapters (FAs). The test environment utilized EMC PowerPath/VE to automatically configure multiple paths, as workloads shifted for the ESX servers and VMs.

The image below clearly shows the FC switch ports for both the ESX servers and the Symmetrix VMAX FA ports (beginning with FA). Each port had a maximum speed of 4 GB.

Interface	Description	VSAN(s)	Mode	Connected To	Speed (Gb)	Rx	Tx
fc8/2	fa 12H ec		11 F	0x0500a7, 50:00:09:72:08:1e:7d:ec	4	42	8
fc4/20	FA 6G 94		11 F	0x0500bd, 50:00:09:72:08:1e:7d:94	4	42	7
fc8/14	fa 9E 21		11 F	0x0500ac, 50:00:09:72:08:1e:7d:21	4	42	7
fc8/8	fa 10H e4		11 F	0x0500a3, 50:00:09:72:08:1e:7d:e4	4	42	8
fc9/7	fa 11G a8		11 F	0x0500a4, 50:00:09:72:08:1e:7d:a8	4	41	7
fc9/20	fa 11E 29		11 F	0x0500a6, 50:00:09:72:08:1e:7d:29	4	40	7
fc7/13	FA 8E 1C		11 F	0x0500bb, 50:00:09:72:08:1e:7d:1c	4	42	8
fc8/7	fa 9G a1		11 F	0x0500ad, 50:00:09:72:08:1e:7d:a1	4	42	7
fc8/16	fa 10E 25		11 F	0x0500a1, 50:00:09:72:08:1e:7d:25	4	42	7
fc4/2	FA 6E 14		11 F	0x0500ba, 50:00:09:72:08:1e:7d:14	4	41	7
fc7/14	FA 8G 9C		11 F	0x0500bc, 50:00:09:72:08:1e:7d:9c	4	41	7
fc9/3	fa 12E 2d		11 F	0x0500a2, 50:00:09:72:08:1e:7d:2d	4	41	7
fc4/13	R900DMX_7_HBA2_P0		11 F	0x050097, 21:00:00:1b:32:8f:6f:31	4	8	42
fc8/23	R900DMX_7_HBA1_P0		11 F	0x05009d, 21:01:00:1b:32:af:6f:31	4	8	42
fc7/23	R900DMX_9_HBA1_P0		11 F	0x050098, 21:00:00:1b:32:8f:e8:30	4	7	41
fc9/5	R900DMX_10_HBA1_P0		11 F	0x0500b5, 21:01:00:1b:32:af:cc:2f	4	7	42
fc7/24	R900DMX_9_HBA2_P0		11 F	0x05009b, 21:00:00:1b:32:8f:2a:31	4	8	42
fc9/4	R900DMX_10_HBA0_P1		11 F	0x050096, 21:00:00:1b:32:8f:85:2e	4	7	41
fc4/16	R900DMX_11_HBA1_P0		11 F	0x0500b6, 21:00:00:1b:32:81:28:9b	4	7	41
fc8/17	R900DMX_8_HBA2_P0		11 F	0x050099, 21:00:00:1b:32:8f:6d:32	4	7	41
fc8/24	R900DMX_11_HBA1_P0		11 F	0x0500b9, 21:00:00:1b:32:81:80:9f	4	7	42
fc8/5	R900DMX_8_HBA1_P0		11 F	0x05009c, 21:00:00:1b:32:82:8e:5e	4	7	41
fc4/6	R900DMX_12_HBA1_P0		11 F	0x0500b8, 21:00:00:1b:32:8f:fa:30	4	7	41
fc4/24	R900DMX_12_HBA1_P0		11 F	0x0500b7, 21:00:00:1b:32:8f:f2:2f	4	7	41

For detailed information on the PowerPath/VE product, visit:

<http://www.emc.com/products/detail/software/powerpath-multipathing.htm>

---

**Symmetrix  
VMAX I/O  
throughput**

The chart below illustrates the total I/O throughput for all 24 Exchange VM Mailbox servers and 120,000 very heavy, .18 IOPS users from the Symmetrix VMAX array perspective.

It shows the array handling about 38,000 Exchange 2010 IOPS during testing. This translates to 1,583 total IOPS generated from each Exchange VM server. The I/O is a combination of random 32k reads and writes for user activity and 256k reads for BDM activity.



---

**Database backup and recovery performance overview**

Two Jetstress tests were performed with a single ESX server and 20,000 users to measure storage backup performance, as follows:

- **Database Read-only Performance**—Measures the sequential read rate of the database files.
- **Transaction Log Recovery/Replay Performance**—Measures the recovery/replay performance by playing transaction logs into the database.

**Note**

DB recovery and Read-only performance tests results were the same for both configurations and are combined in a single table for easy review; see [Testing methodology and results>Backup VSS test-database Read-only performance](#).

---

**Backup VSS test: Database Read-only performance**

The backup VSS Jetstress test measures the maximum rate at which databases can be backed up using VSS. The following table shows test results for the average rate for a single database file.

<b>MB read/s per database</b>	83.24
<b>MB read/s total per server</b>	665.92

---

**Soft Recovery test: Transaction log recovery/replay performance**

The Soft Recovery Jetstress test measures the maximum rate at which the log files can be played against the databases. The following table shows test results for the average rate for 500 log files played in a single database. Each log file is 1 MB in size.

<b>Average time to play one log file</b>	1.4 sec
--	---------

---

---

## Validating the performance impact of thin versus regular devices with Jetstress

---

### Introduction

This section compares thin versus regular devices in the Exchange 2010 test environment. Regular devices can also be referred to as “thick” devices and represent fully-provisioned, traditional LUNs. This white paper uses the term “regular” devices throughout.

---

### Virtual Provisioning testing approach

We wanted to answer one essential question during our initial testing of using thin devices in the Exchange 2010 test environment:

#### **What is the performance impact of using thin devices for Exchange 2010 database and log volumes, versus using regular devices?**

This question was answered by taking four mailbox server VMs (20,000 users housed on a single ESX server) and converting their storage from regular devices to data devices. The data devices were added to a thin pool (one pool per VM). Each thin pool was populated with data devices created from its original 16 spindle disk group.

Jetstress was run against a single ESX server containing four VMs and 20,000 users whose database and logs were on thin devices. Jetstress was also run against a second ESX server containing 20,000 users but with regular devices for their database and log volumes at the same time. All mailbox VMs applied the same Jetstress profile:

.18 IOPS with a 1 GB Mailbox quota

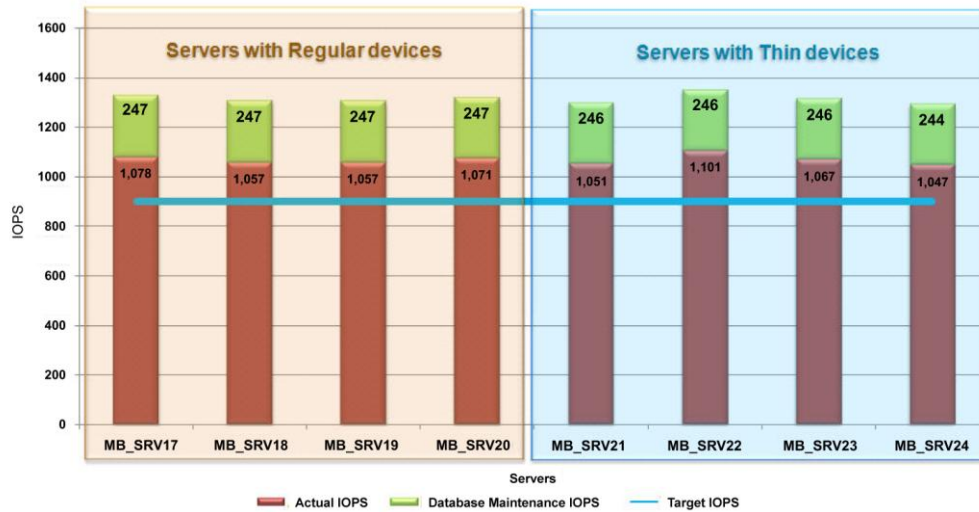
---

**Test results—  
Comparing thin  
devices versus  
regular devices**

Comparatively speaking, the chart below shows that the thin devices perform equal to the regular devices for Exchange 2010 I/O:

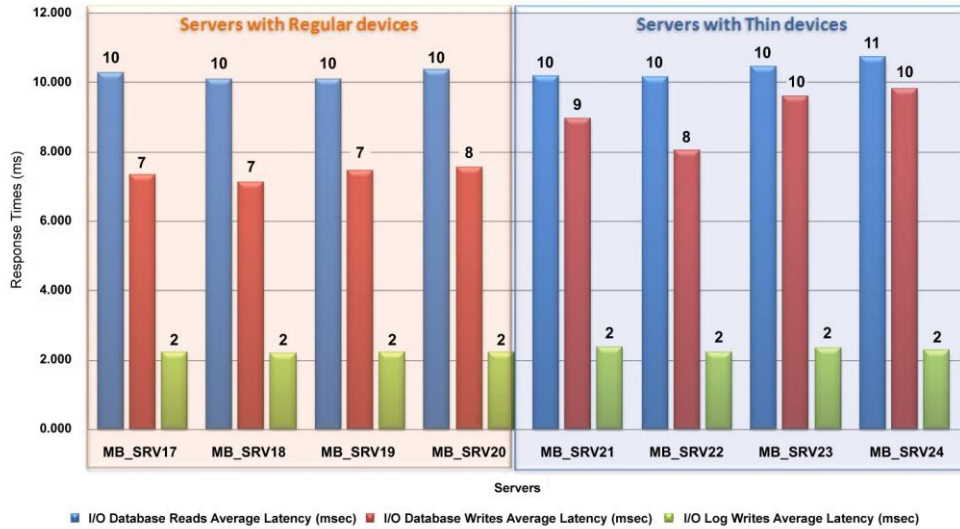
- Thin devices produced 4,266 Exchange IOPS
- Regular devices produced 4,263 Exchange IOPS

From an Exchange 2010 I/O perspective, there is no impact putting your database and logs on thin devices versus regular devices.



**Test results—  
Comparing thin  
device  
latencies  
versus regular  
device  
latencies**

When comparing the Jetstress 2010 database and log latency results a small 2 ms difference was observed on the database write latency. All other latencies were the same. Overall Symmetrix VMAX thin devices and Virtual Provisioning are an excellent platform for Exchange 2010 to live, providing unsurpassed flexibility for Exchange administrators.



---

## vSphere mailbox server validation testing with LoadGen

---

<b>Introduction</b>	This section details the ESX server and mailbox VM validation testing with Microsoft's LoadGen tool.
<b>Microsoft LoadGen</b>	LoadGen is a simulation tool to measure the impact of MAPI, OWA, IMAP, POP and SMTP clients on Exchange servers. LoadGen is used to validate the e-mail infrastructure.
<b>LoadGen testing approach</b>	LoadGen is used to validate the vSphere Exchange 2010 mailbox building-block's performance under load. In addition, Exchange 2010, EMC storage, and VMware performance data were logged for analysis during an 8-hour test run.
<b>LoadGen test configuration</b>	<p>LoadGen test results are provided based on the following configuration:</p> <ul style="list-style-type: none"><li>• 20,000 very heavy Exchange 2010 users (8-hour duration)</li><li>• Single ESX server across 4 VMs</li></ul> <p>This test is used to validate that the Exchange 2010 mailbox VM building-block methodology was sound and performed consistently under user mail load.</p>
<b>LoadGen mail client performance results</b>	<p>The LoadGen very heavy profile produced approximately 13 log files per user for a total of approximately 260,000 logs during the 8-hour test run. In an effort to stress the environment as much as possible, the LoadGen test was run in Outlook 2007 Online mode. The LoadGen profile executed 178 tasks per user for a total of 3,560,000 tasks during a test run.</p> <p>Exchange client performance is always a critical factor. To gauge the client experience during testing, Exchange performance monitor counters were collected on each of the four mailbox servers. The Exchange user response times and latency were excellent, during LoadGen testing:</p> <ul style="list-style-type: none"><li>• Average RPC latency—Approximately 2.6 ms (Microsoft recommends that average RPC latency does not exceed 10 ms.)</li><li>• RPC requests—Averaged three requests with a maximum of 17 (keep this counter below 70 at all times.)</li></ul>

---

**LoadGen VM  
CPU utilization  
results**

Two of the primary performance counters that need to be examined from a vSphere perspective include ESX Server and VM CPU utilization. The following chart represents the CPU utilization on the Exchange ESX server with four mailbox VMs during the 8-hour LoadGen test.

LoadGen ran at a very heavy load for 8 hours. As can be seen from the ESX **Performance** tab, the mailbox VMs averaged around 53 percent CPU utilization with spikes up to 75 percent while LoadGen was running. The ESX servers averaged around 55 percent CPU utilization. These results validate the design, and show a well-performing ESX server and Exchange 2010 mailbox VMs.

**Note**

CPU sizing is based on guidance from Microsoft TechNet:  
<http://technet.microsoft.com/en-us/library/ee712771.aspx>



---

## Conclusion

---

### Summary

The EMC Symmetrix VMAX storage array, combined with VMware vSphere, the industry's first cloud operating system, are leveraged to consolidate and virtualize the Exchange 2010 application. EMC delivers the power of EMC Symmetrix VMAX and VMware vSphere advanced virtualization to maintain efficient use of resources, HA, and database protection.

This solution demonstrates these capabilities through the testing of 120,000 very heavy Exchange 2010 users with a 1 GB mailbox quota. By consolidating Exchange servers with a virtualized infrastructure, organizations can achieve significant cost reductions in the areas of network management and resource consolidation.

The methodologies used to test enterprise Exchange Server 2010 on the Symmetrix VMAX storage array will help organizations to deploy Exchange 2010 more easily and efficiently.

---

### Findings

The following list details the findings of this white paper.

- Symmetrix Virtual Provisioning can lead to significant cost savings at initial deployment, particularly when the ultimate goal is to provide large Exchange 2010 mailboxes for all users with minimal deployment costs.
  - Thin devices on a Symmetrix VMAX perform as well as regular devices when designed properly.
  - A reduction in physical servers for Exchange mailboxes is achieved (from 24 to 6 vSphere ESX servers).
-

---

## References

---

### White papers

For additional information, see the white papers listed below (available on EMC Powerlink®):

- *Implementing EMC Symmetrix Virtual Provisioning with VMware vSphere—Applied Technology*
  - *EMC Symmetrix VMAX and Microsoft Exchange Server—Applied Technology*
  - *EMC Symmetrix VMAX and VMware Virtual Infrastructure—Applied Technology*
  - *Secure and Consolidated 16,000 Exchange Users Solution on a VMware/EMC Environment—Applied Technology*
- 

### Product documentation

For additional information, see the product documents listed below.

- *EMC Symmetrix Management Console (SMC) online help* (integrated with the Symmetrix VMAX)
- 

### Other documentation

See the following VMware technical documents for further reference:

- VMware Certified Compatibility Guides
-