

EMC Virtualized Architecture for Microsoft Exchange Server 2007 with VMware Virtual Infrastructure 3 and EMC CLARiiON CX4-960

Best Practices Planning

Abstract

This white paper presents best practices and validated test results for building a Virtualized Microsoft Exchange Server 2007 environment using VMware ESX Server and EMC® CLARiiON® CX4 storage. The design enables customers considering an Exchange 2007 deployment to capitalize on the benefits of a virtualized platform.

July 2009

Copyright © 2009 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 software 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 H6121

Table of Contents

Executive summary	5
Introduction	5
Audience	5
Terminology	6
Overview	6
EMC CLARiiON CX4-960 storage	7
Features	7
Benefits	8
VMware ESX Server	8
Benefits of a virtualized Microsoft Exchange Server 2007 deployment	9
Readiness Assessment for Virtualized Microsoft Applications	10
Environment configuration best practices for virtualizing Microsoft Exchange Server 2007	10
ESX Server configuration best practices	10
Guest VM configuration best practices	11
Sizing Exchange server VM best practices.....	11
Exchange VM questions for planning	12
Exchange storage questions for planning	13
EMC CLARiiON CX4-960 storage array design considerations.....	14
Exchange 2007 Mailbox Server building-block design overview	14
What is a building-block?	14
How the building-block approach is implemented.....	15
Creating a building-block.....	16
Step 1. Identify user requirements	16
Step 2. Identify and calculate storage requirements	16
Step 3. Identify Exchange VM requirements	17
Step 4. Finalize and define your Exchange VM building block	17
Validation of the virtualized Exchange 2007 environment.....	17
Testing methodologies.....	17
Exchange performance validation tools	17
Performance collection and data analysis tools	18
CLARiiON CX4-960 performance validation with JetStress based on the Microsoft ESRP testing framework	18
Microsoft ESRP testing overview.....	18
JetStress testing—phase 1 summary	19
Simulated Exchange configuration for JetStress testing phase 1	19
Primary storage hardware	21
Primary storage software.....	21
Primary storage disk configuration (mailbox store disks).....	22
Primary storage disk configuration (transactional log disks).....	22
Performance testing across server metrics.....	23

Reliability	23
Storage performance results.....	23
Individual server metrics	24
Aggregate performance across all servers metrics	25
JetStress testing—phase 2 summary	25
Simulated Exchange configuration for JetStress testing phase 2	26
CLARiiON storage hardware used in this configuration.....	28
Disk configurations.....	28
Individual server metrics	29
Aggregate performance across all server metrics	30
Streaming backup/recovery performance	31
Database read-only performance test.....	31
Log read-only performance test	31
CLARiiON CX4-960 performance	32
JetStress testing—phase 3 summary	33
JetStress testing Phase 3—CLARiiON storage processor simulated failure test.....	33
Exchange 2007 environment validation with LoadGen	34
Why is LoadGen testing necessary?	34
Measuring good Exchange performance	34
LoadGen testing—environment configuration.....	35
Virtualized Exchange environmental performance	37
Exchange HUB/CAS performance	40
Conclusion.....	40
References	41

Executive summary

In recent years the features of server virtualization platforms have become much more robust, and the significant benefits attainable from a virtualized infrastructure have become clearer. Spurred by increasing requirements for high availability and cost efficiency, organizations have been consolidating underutilized servers in their data centers into virtualized environments at a remarkable pace. More importantly, critical applications like Microsoft Exchange Server can now take full advantage of virtualization technology to provide more efficient use of physical infrastructure and 64-bit architecture.

An EMC virtualized architecture supporting a Microsoft Exchange Server 2007 environment enables customers to leverage consolidated hardware technology, and best use of physical assets such as storage, network, and computing resources to control costs and respond quickly to changing business needs. The virtualization approach to IT management creates virtual services out of the physical infrastructure, enabling administrators to allocate these virtual resources quickly to the business units that need them most.

Introduction

This white paper presents methodologies and best practices, based on validated test results for designing a virtualized Microsoft Exchange Server 2007 environment using VMware Infrastructure 3, and EMC[®] CLARiiON[®] storage. Specifically, sizing Exchange virtual machines (VMs) using a building-block approach is discussed in the paper. The building-block methodology will help Exchange administrators deploy Exchange more easily and efficiently.

In addition, as part of the performance validation process, the EMC CLARiiON CX4-960 array was configured to support up to 44,000 users with various profile types and mailbox sizes up to 2 GB. Further, this document presents CLARiiON CX4-960 storage performance data using the same guidelines formulated by the Microsoft Exchange Solution Reviewed Program (ESRP) for storage vendors.

EMC has devoted a great deal of effort to testing integrated solutions for Microsoft Exchange with VMware on CLARiiON storage. A complete list of supporting documentation is provided in the "[References](#)" section at the end of this document.

Audience

This white paper is intended for:

- Customers, including IT planners, storage architects, and administrators
- EMC technical staff and partners
- Field personnel, who are tasked with implementing Microsoft Exchange Server 2007 solutions involving EMC CLARiiON CX4 storage in a virtualized environment
- Engineering and product development groups
- Account personnel involved in pre-sales activities

This white paper assumes the reader has a basic knowledge and understanding of VMware Infrastructure, Microsoft Exchange Server 2007, Microsoft Windows, and the CLARiiON storage array.

Terminology

This white paper references the following terms.

Building-block—Represents the required amount of resources required to support a specific number of Exchange 2007 users on a single VM. The amount of required resources is derived from a specific user profile type.

Distributed Resource Scheduler (DRS)—A VMware feature to ensure maximum performance and load balancing across an ESX cluster in a reported or automated fashion. VMotion is employed to migrate candidate VMs.

EMC Replication Manager—Coordinates the entire data replication process—from discovery and configuration—to the management of multiple application consistent disk-based replicas.

Microsoft Exchange Solution Reviewed Program (ESRP)—Standards developed by Microsoft Corporation to provide a common storage testing framework for storage vendors to present Microsoft Exchange Server software test data.

Raw Device Mapping (RDM)—A method to expose a LUN to a VM without hypervisor translation.

Virtual LUN—Enhanced Virtual LUN technology that enables data migration within an array without host or application disruption. Virtual LUN brings a tiered storage strategy to life by easily moving information throughout the storage system as its value changes over time. It can assist in system reconfiguration.

VMotion—A method to migrate an active VM from one ESX host to another with no perceived impact.

VMware HA (High Availability)—VMware software that offers extended VM protection using quick restart on hardware failure.

Overview

When Exchange administrators are faced with the task of implementing a plan for new Exchange 2007 deployment on VMware or a plan for migration of the Exchange 2007 environment from physical to virtual, they need to have a solid understanding of what is necessary to achieve optimal performance in a Virtual Infrastructure environment.

EMC and VMware have performed multiple validation tests and established specific guidelines and best practices for various Exchange 2007 configurations. In this white paper, EMC performed comprehensive testing in order to guide Exchange administrators and provide methodologies, guidelines and best practices to deploy Exchange 2007 on an EMC CLARiiON array and VMware ESX 3.5. In the tested configuration environment, three simple building-blocks were created to validate the design methodology. Different Exchange user profiles—from 0.48 IOPS to 2 IOPS—were tested along with mailbox sizes ranging from 500 MB to 2 GB were used in designing each of the building-blocks.

Each building-block represents an amount of resources that are required to support a specified number of users on a single Exchange 2007 Mailbox Server role. To ensure that the Exchange performance is within recommended Microsoft guidelines, rigorous validation and testing was performed using recommended Exchange performance validation tools such as Microsoft JetStress and Microsoft LoadGen.

EMC CLARiiON CX4-960 storage

The EMC CLARiiON storage array family leads the midrange storage market in providing customers with cost-effective solutions that deliver the highest levels of performance, functionality, and reliability; see [Figure 1](#).

The unique combination of a breakthrough hardware design and advanced software capabilities enables EMC CLARiiON CX4 series systems to meet the growing IT challenges of today's midsize and large enterprises by:

- Scaling system capacity and performance
- Simplifying management in complex environments
- Delivering increasing levels of information availability and protection for mission-critical applications

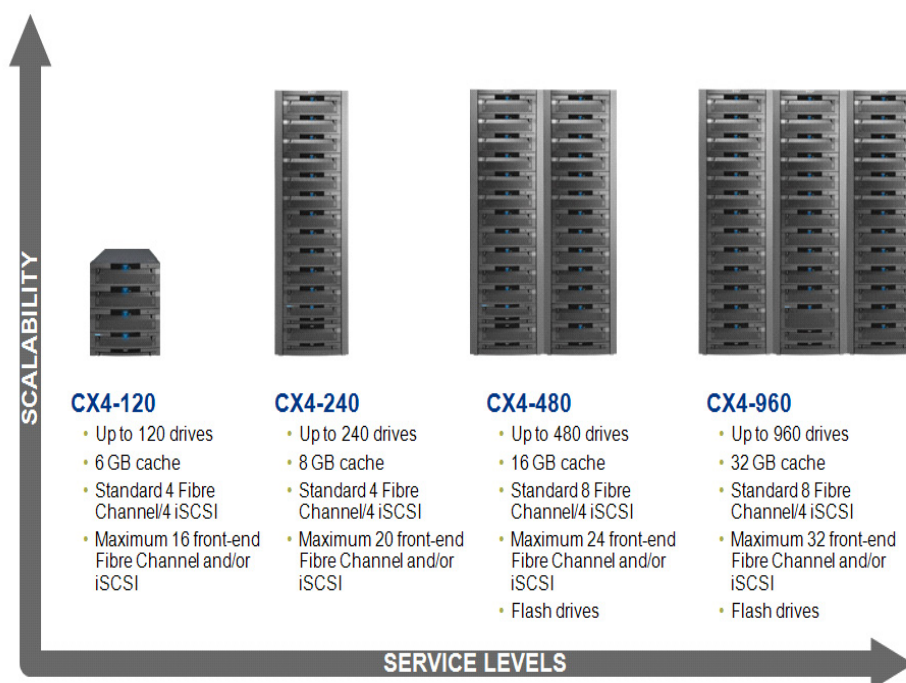


Figure 1. CLARiiON CX4 midtier array products

Features

The EMC CLARiiON CX4-960 delivers performance, scalability and advanced data management features in one, easy-to-use storage solution. The EMC CLARiiON CX4 series integrates industry-changing, midrange storage innovations that include:

- **Flash drives** offering Tier “0” storage with 30 times more IOPS performance than Fibre Channel drives
- **UltraFlex™ technology** for dual protocol connectivity, online expansion via I/O modules, and readiness for future technologies—such as 8 Gb/s Fibre Channel (FC) and 10 Gb/s iSCSI
- **CLARiiON Virtual Provisioning™** for increased utilization, just-in-time capacity allocation, and simplified provisioning and capacity planning
- **Multi-core Intel Xeon processors, increased memory, and 64-bit FLARE®**, providing up to twice the performance and scale, and the foundation for more advanced software functionality
- **Low power SATA II drives, adaptive cooling, and disk drive spin down** providing new energy-efficiency options for customers

-
- **Concurrent local and remote replication (CLR)** through the integrated EMC RecoverPoint splitter that simplifies CLARiiON CX4 series deployments and allows for asynchronous local and remote replication for recovery to any point in time, plus WAN bandwidth reduction.

Benefits

CLARiiON CX4-960 provides these benefits:

- **Lowest cost of ownership.** The CX4-960 consolidates twice the workload in a single system. Compared to the previous generations of CLARiiON products, the CX4 architecture delivers twice the capacity scale (up to 960 drives), twice the memory, and twice the LUNs. Competing storage providers make customers deploy multiple arrays to equal the performance and scalability of a single EMC CLARiiON CX4.
- **New energy efficiency.** EMC CLARiiON CX4-960 provides new energy efficiency through support for Flash drive technology providing lower energy requirements than traditional drives.
- **Site-level protection.** CLR replication with EMC RecoverPoint provides maximum protection for physical and virtual environments.
- **Optimized for virtualized environments.** The VMware-certified CLARiiON CX4 series features the following functionality:
 - **Shared storage.** Gives customers a shared storage solution that is integrated with VMware VMotion for maximum flexibility.
 - **RecoverPoint and VMware Site Recovery Manager (SRM).** Customers can leverage VMware SRM using RecoverPoint or MirrorView™/S software to simplify and automate disaster recovery (DR) setup, failover, failback, and testing, while ensuring the proper execution and management of recovery plans.
 - **Resource management.** Integrated Navisphere® Quality of Service (QoS) Manager with VMware Distributed Resource Scheduler (DRS) provides end-to-end application service-level management from the virtual server to CLARiiON storage.

VMware ESX Server

VMware ESX Server consists of virtualization software that provides server consolidation by allowing several instances of similar and dissimilar operating systems to run as VMs on one physical machine. This cost-effective, highly scalable VM platform offers advanced resource management capabilities. VMware ESX Server minimizes the total cost of ownership (TCO) of computing infrastructure by:

- Increasing resource utilization
- Decreasing the number of servers and all associated costs
- Maximizing server manageability

ESX Server runs directly on the hardware, allowing VMs to run on top of the virtualization layer provided by VMware. The combination of an operating system and applications is referred to as a virtual machine, or VM. VMware vCenter is the VMware management interface used to administer one or more ESX Servers.

Benefits of a virtualized Microsoft Exchange Server 2007 deployment

Many organizations have been transitioning their Exchange environments to the VMware virtual platform. Consider the potential benefits by virtualizing a Microsoft Exchange environment with VMware and CLARiiON storage:

- **Best utilization of hardware resources that translates to reduced costs**—Virtualization can help companies lower costs associated with their Exchange server infrastructure by consolidating Exchange servers into VMs running on powerful 64-bit server hardware. These cost savings include server hardware and maintenance costs, reduced rack space and floor space, reduced data center power and cooling costs, and reduced infrastructure costs (host bus adapters, cables and switch ports) among others.
- **Increased flexibility with Exchange 2007 roles**—VMs run independently of the underlying physical server hardware, providing substantial improvements in overall flexibility. Virtualized Exchange servers are no longer tied to a physical server and can be dynamically load balanced across VMware ESX Servers with VMware Distributed Resource Scheduler (DRS).
- **Unprecedented control of Exchange servers with VMware VMotion**—Virtualized Exchange servers can be manually moved live across VMware ESX Servers at any time with VMware VMotion, allowing the Exchange environment to rapidly adjust to changing requirements.
- **Simplified Exchange Mailbox Server design**—Using the building-block approach described in this document, Exchange environments can be deployed using a modular approach. Each building-block in this design has been pre-sized and tested to ensure optimal performance. Using the building-block approach takes the guesswork out of sizing Mailbox Servers and associated storage. The building-block approach allows the environment to scale easily using a predictable approach, where each building-block has a well-understood performance profile.
- **Improved service-level availability and performance**—Running Exchange 2007 on a VMware virtual infrastructure allows an organization to improve service level agreements (SLA) and reduce the risk of downtime associated with planned and unplanned outages.
- **VMware HA with less complexity**—With any critical application like Exchange, there must be a mechanism for the application to quickly recover from the hardware failure. In the case of an unplanned hardware outage or unexpected guest OS failure, VMware HA can automatically restart the VM and DRS will load balance the Exchange VMs on the remaining physical ESX Servers within the cluster.
- **Microsoft support considerations for VMware**—Microsoft officially supports VMware ESX for running Microsoft Windows and major applications including Microsoft Exchange Server 2007, Microsoft SQL Server 2008, and Microsoft Office SharePoint Server 2007. Customers can now run Exchange on VMware with peace of mind knowing that support is available. For more information on Microsoft's Server Virtualization Validation Program visit:
<http://windowsservercatalog.com/svvp.aspx?svvppage=svvp.hm>
- **Unprecedented server scalability and return on investment (ROI)**—Today's new 64-bit servers required for Exchange 2007 continue to ship with multi-core processors and increased memory density. Exchange 2007 does not scale well in an un-virtualized deployment beyond more than two sockets, eight processor cores, and 32 GB RAM.

Readiness Assessment for Virtualized Microsoft Applications

A new EMC service called Readiness Assessment for Virtualized Microsoft Applications is available to provide analysis of the current environment. This service involves conducting a thorough evaluation of your existing configuration and requirements, and delivers detailed recommendations on how to optimize the deployment of your new virtualized configuration. For more details visit the EMC Powerlink[®] website.

Environment configuration best practices for virtualizing Microsoft Exchange Server 2007

As with any Exchange environment, it's important to build high availability into the entire virtualized infrastructure. This involves avoiding single points of failure, and planning for failover timings and handling data loss at levels that meet your requirements. The major components involved in the configuration of this solution include Microsoft Exchange 2007 virtualized on a VMware ESX 3.5 platform with EMC CLARiiON CX4-960 storage.

Use the following configuration guidelines and best practices as a reference point to begin planning your deployment:

- **Be sure to provide high availability for the active directory environment and DNS.** Domain Controllers and DNS—Windows domain controllers (DCs) with DNS services run well on VMs. In general, for every four Exchange servers, there should be one DC/GC server. For fault tolerance you should have at least two DCs with global catalog (GC) enabled as well as at least two DNS servers per site in order to minimize replication, improve lookup times, and reduce network traffic.
- **Avoid using DHCP for critical servers.** DHCP is often a good option for assigning IP addresses for client machines. However, for best availability it is best practice to use static IP addresses for Exchange and other key servers and have them registered in DNS.
- **Keep the infrastructure time in synch.** To make administration and troubleshooting easier, make an effort to keep the times synchronized among the various infrastructure components (including domain controllers, storage arrays, ESX Servers and their VMs).
- **Configure each ESX Server with at least two HBAs for FC connections or high speed network interface cards (NICs) for iSCSI connections.** For highest resiliency these should represent independent paths, with each going to a different fabric or network switch. The corresponding two connections from each switch to the array should also connect to separate storage processors (SPs). With iSCSI, use two dedicated gigabit iSCSI network paths to the CLARiiON storage system.

ESX Server configuration best practices

Because ESX Servers are usually used to support the equivalent function of multiple standard servers, they are typically configured more robustly. Best practices include:

- ESX Virtual Infrastructure version 3.5 or later
- Number of CPU cores equal to the number of virtual CPUs you plan to run concurrently (within the set of VMs on the ESX Server)
- Physical memory equivalent to the sum of the memory recommendations by Microsoft for the individual Exchange server VMs, plus additional memory required for the ESX Server
- Configuring each ESX Server with at least two HBAs for FC connections or high speed network interface cards (NICs) for iSCSI connections.

For highest resiliency these should represent independent paths, with each going to a different fabric or network switch. The corresponding two connections from each switch to the array should also connect to separate SPs. With iSCSI, use two dedicated gigabit iSCSI network paths to the CLARiiON storage system.

Guest VM configuration best practices

With VMware Virtual Infrastructure 3, ESX Server 3.5 guest VMs can be configured with up to four vCPUs and up to 64 GB memory. This is usually more than enough to support a Mailbox Server with up to 4,000 very heavy Exchange 2007 users.

Best practices for configuring VMs for an Exchange environment include:

- **Always install VMware tools on guest VMs.** Installing VMware tools includes an optimized network driver, provides enhanced user experiences, and installs memory reclamation drivers. Refer to <http://kb.vmware.com/kb/340> for more information.
- **Disable unused virtual devices.** Disconnect the unused CD connected to the ISO files and devices such as COM ports, LPT ports, floppy drives, and USB adapters are periodically polled by the guest operating system even if they are not in use. To avoid this unproductive polling waste, we recommend disabling those devices in VM configurations. Also, it is important to disconnect unused CD connected to ISOs.
- **Place VM swap space on fast storage.** ESX Server creates a swap file per VM that is equal in size to the difference between the VMs, configured memory size and its reservation. Make sure that the VM swap file is placed on a high-speed storage system. The default is to place it on the same datastore as the guest operating system. In this case be sure to allocate the added capacity as necessary.
- **Set the memory reservation:** Set the memory reservation to the configured size of the VM, which results in a per-VM VMkernel swap file of zero bytes. The guest operating system within the VM will maintain its own separate swap/page file.
- **Windows page file configuration.** Include a page file on both boot partition and on additional drives; refer to <http://support.microsoft.com/kb/197379/en-us>. To find the appropriate page file size for 64-bit Windows 2003, refer to <http://support.microsoft.com/kb/889654/en-us>.
- **Increase service console memory.** If planning to install hardware agents or backup agents on the ESX Server, increase the console memory from 272 MB to a higher number depending on need.
- **Rename the ESX datastore for easy identification.** Change the default datastore naming to something that can be easily identified.

Sizing Exchange server VM best practices

Planning can be started for designing an Exchange VM based on specific user profile requirements. First, provide answers to the following questions:

- ["Exchange VM questions for planning" on page 12](#)
- ["Exchange storage questions for planning" on page 13](#)

Exchange VM questions for planning

Answer the following questions:

- **How many users can we place on an Exchange Mailbox Server VM?**
 - Approximately 500 to 700 Exchange IOPS per core (based on EMC validation testing performed)
 - 1,000 very heavy (0.48 IOPS) users per VM core (4,000 users per VM with four vCPUs)
 - 375 @ 2 IOPS users per VM core (1,500 users per VM with four vCPUs)

Follow Microsoft's guidelines for Mailbox Server Storage Design; see [Table 1](#):

Note: The user profile types listed below are based on the Microsoft *Mailbox Server Storage Design* article available at <http://technet.microsoft.com/en-us/library/bb738147.aspx>.

Table 1. Exchange 2007 user profile types

User type	Send/receive per day (50 k messages)	Database cache / user	Estimated IOPS/user
Light	5 sent / 20 received	2 MB	0.11
Average	10 sent / 40 received	3.5 MB	0.18
Heavy	20 sent / 80 received	5 MB	0.32
Very heavy	30 sent / 120 received	5 MB	0.48
Extra heavy	40 sent / 160 received	5MB	0.64

- **How many Exchange VMs are required per ESX Server?**
 - Depends on the server type, number and speed of CPU cores, and memory requirements.
 - In the tested configuration, the physical ESX Server with four quad-core CPUs was able to support four Exchange VMs.
- **How much memory is needed per VM and ESX Server?**

This will depend on the number of users per Exchange VM.

 - Use Microsoft best practice of 5 MB per user, plus 2 GB
 - Need additional memory based on ESX Server memory overhead best practices
- **What type of VMware volumes should I use?**
 - Use the VMFS volume type for the Exchange server guest OS. This volume should be large enough to accommodate the OS and page file.
 - Use RDMS or VMFS as the volume type for Exchange database and log volumes. Sometimes RDMS will provide better granularity and integration with other EMC and third-party backup software.

Exchange storage questions for planning

Answer the following questions:

- **How many drives are needed per Exchange server, and VM?**
- Calculate database spindle requirements:

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

Where:

IOPS = The number of I/Os per second

%R = Percent of read I/Os

%W = Percent of write I/Os

WP = RAID write penalty multiplier (RAID 1=2, RAID 5=4)

Physical Disk Speed = 180 for 15k rpm drive or 140 for 10k rpm drive

- Calculate transaction logs spindle requirements:

$$\frac{25\% \text{ of DB IOPS}}{\text{Physical Disk Speed}} = \text{Required Physical Disks}$$

- **What RAID type should I use (this is dependent on I/O and space requirements)?**

- RAID 1/0 is recommended for best performance
- For configurations with larger mailboxes RAID 5 often makes sense for database LUNs
- RAID 5 is recommended for VMs, OS datastores

- **How many ESGs per server?**

- Depends on a number of issues, including backup and restore requirements, and online maintenance (OLM) performance
- A 200 GB Exchange database worked well in both the SAN and replication environments
- May need to add multiple ESGs per LUN (due to the ESX 3.5 LUN limit of 256 and 60 LUNs per VM)

- **How many storage groups per database/log LUN pair?**

In most cases, two LUNs should be allocated for each ESG—one for the transaction logs, and one for the database file. Rapid recovery VSS-compliant solutions typically recover an entire LUN at a time. By placing each ESG on its own LUN pair, recovery is faster and affects only users in that storage group.

However, especially on an Exchange server with a large number of storage groups or server that is virtualized with ESX 3.5, there are cases where it is reasonable to use the same database/log LUN pair to hold more than one storage group because:

- If the backup/recovery method chosen is at a file level rather than a LUN level, it's reasonable to configure multiple storage groups each in their own folders on the same LUN. As with all other cases of data placement, always consider your backup and replication schedules, and combine storage groups on LUNs that result in balanced I/O across the physical drives.
- By residing on the same LUN, multiple ESGs can share the same free space for growth, or offline defragmentation.

EMC CLARiiON CX4-960 storage array design considerations

For best performance, it is recommended to follow EMC best practices when configuring the CLARiiON storage array; see the *EMC CLARiiON Best Practices for Performance and Availability* white paper available on the EMC Powerlink website.

General CLARiiON configuration best practices for Exchange include:

- Enable read and write cache on the array. Assign at least 80 percent of cache to write cache and the remaining amount to the read cache.
- Keep all other CLARiiON default configuration settings. CLARiiON tuning parameters are set with default values designed to provide optimized performance for the entire array. In general their values should not be changed:
 - The default CLARiiON cache page size of 8 KB matches well with the 8 KB page size of Exchange 2007
 - High/low watermarks for memory cache management
 - Cache settings on individual LUNs (“On” by default for fibre and SATA; “Off” by default for EFD)
- Take best advantage of CLARiiON performance and availability features by balancing the Exchange load as well as possible across the array storage processors, system buses, available set of drives, and in larger environments—across arrays.
- To ensure the target ID of LUNs are the same across the ESX hosts in a cluster and in the HA configuration, assign all the ESX Servers to the same CLARiiON storage group.
- Install Navisphere Host Agent and CLI on the ESX Server for automatic host registrations on CLARiiON and ease of management.
- Install the supported HBA hardware and driver on the ESX Server.
- Use the Microsoft diskpart utility (in Windows 2003 SP1) to align all Microsoft-Exchange-related disks, using a value of 64 for CLARiiON. This aligns all of the Exchange-related NTFS partitions on a 64-KB boundary. This is not necessary if Windows 2008 is used as the base OS.
- In Windows 2003 or 2008, format Exchange-related volumes use a 64 k allocation unit size.
- Separate Exchange logs and databases onto different disks and CLARiiON RAID groups.
- Size and configure the environment for spindle performance as a primary consideration, with spindle or storage capacity a secondary issue. In other words, size for performance first and make the capacity requirements work.

Exchange 2007 Mailbox Server building-block design overview

This deployment implemented the Exchange Mailbox Server VM building-block approach, adding flexibility and ease of use for Exchange and storage administrators. Now, armed with all of the information provided earlier in this white paper, we can create Exchange VMs building-blocks.

What is a building-block?

A building-block represents the required amount of resources required to support a specific number of Exchange 2007 users on a single VM. The amount of required resources is derived from a specific user profile type, mailbox size and disk requirements.

Using the building-block approach removes the guesswork and simplifies implementation of Exchange VMs. Once the initial Exchange Mailbox Server VM 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 additions much easier and straightforward.

How the building-block approach is implemented

To assist with your Exchange 2007 VMs deployments, we have created and validated three different building-blocks that are based on various Exchange user profile types and mailbox sizes. Understanding how these building-blocks are created will help you to configure your own, based on your Exchange user profile requirements.

The three different building-blocks that we have created and validated for performance include (each of these building-blocks were configured based on the user requirements):

- **Building-block 1**—Based on 4,000 very heavy (0.48 IOPS) Exchange 2007 users with 500 MB mailboxes. Each of the 4,000 users occupied a single Exchange 2007 Mailbox Server VM, and the users were divided among 10 ESGs, with 400 users per each storage group.
- **Building-block 2**—Based on 4,000 very heavy (0.48 IOPS) Exchange 2007 users with 2 GB mailboxes. Each of the 4,000 users occupied a single Exchange 2007 Mailbox Server VM, and users were divided among 40 ESGs, with 100 users per each storage group.
- **Building-block 3**—Based on 1,500 Exchange 2007 users whose data load exceeds the extra heavy profile (for example, 2 IOPS, with 2 GB mailboxes). Each of the 1,500 users occupied a single Exchange 2007 Mailbox Server VM, and users were divided among 15 ESGs, with 100 users per each storage group.

Additional details for each of the tested building-blocks is provided; see [Table 2](#) and [Table 3](#).

Table 2. Exchange VM building-block details

Parameters	Building-block 1	Building-block 2	Building-block 3
Number of users per Exchange Mailbox Server/VM	4,000	4,000	1,500
User IOPS profile	Very heavy (0.48 IOPS)	Very heavy (0.48 IOPS) spindle requirements based on space	Custom (2.0 IOPS)
Mailbox size	500 MB	2 GB	2 GB
vCPUs per VM	Four	Four	Four
Memory per VM	24 GB	24 GB	12 GB
Disk configuration	20 disks RAID 10 for DB and Logs (16 450 GB 15k for DB 4 146 GB 15k for Logs)	39 disks RAID 5 for DBs (35 450 GB 15k for DB 4 450 GB 15k for Logs)	28 disks RAID 10 for DB and Logs (24 450 GB 15k for DB 4 450 GB 15k for Logs)
ESG layout	10 (400 user per ESG) 2 ESGs per LUN (log LUNs same way)	40 (100 users per ESG) 2 ESGs per LUN (log LUNs same way)	15 (100 users per ESG) 2 ESGs per LUN (log LUNs same way)

Based on current ESX 3.5 storage limitations, and to ensure that each ESX Server can handle our building-blocks and have room for future growth, the LUNS were created as described in [Table 3](#).

Table 3. CLARiiON LUN configurations per Exchange VM building-block

Parameters	Building-block 1	Building-block 2	Building-block 3
CLARiiON/ESX/VM LUN configurations	12 LUNs per VM - 1 OS VMFS - 1 mount point root -10 Exchange RDMs	42 LUNs per VM - 1 OS VMFS - 1 mount point root - 40 Exchange RDMs	20 LUNs per VM - 1 OS VMFS - 1 mount point root - 15 Exchange RDMs

Creating a building-block

Let us walk through the process of how a building-block is created. For this exercise we will use the user profile described in building-block 1. The process of creating a building-block involves four simple steps:

1. Identify user requirements.
2. Identify and calculate storage requirements.
3. Identify Exchange VM requirements.
4. Finalize and define your Exchange VM building-block

Step 1. Identify user requirements

Need to support 4,000 users with 0.48 IOPS per user and a 500 MB mailbox quota.

Step 2. Identify and calculate storage requirements

Need to support 1,920 IOPS (4,000 * 0.48) and have 450 GB 15k drives for the databases with 146 GB 15k drives for logs on the CLARiiON array.

- Required database spindles = **16** spindles in a RAID 10 configuration:
 $(1,920 * .50) + 2(1,920 * .50) / 180 = 16$
- Required log spindles = **4** spindles in a RAID 10 configuration:
 25% of DB IOPS is $384 / 180 = 4$ (rounded up)
- Total of **20** spindles would be required to support the required number of users and IOPS

How much storage space is needed and how many ESGs do we need to create?

- Follow best practices and keep databases within 200 GB. When all calculations are completed, 2,850 GB (2,600 GB DBs + 250 GB logs) of storage space will be required to support the required number of users.
- Database size: Having 400 users per database will provide us with a 200 GB database size:
 $500 \text{ MB} * 400 \text{ users} = 200 \text{ GB}$
- Database LUN size requirements: Will add in Exchange database overhead requirements:
 $200 \text{ GB} + 30\% \text{ (indexing + white space)} = 260 \text{ GB (2,600 GB per server)}$
- Log LUN size requirements: With 12 logs generated per user per day (based on test results) * 400 users = $4.8 \text{ GB} * 5 \text{ days retention} \sim 25 \text{ GB (250 GB per server)}$
- Taking into consideration VMware storage guidelines, we will place two ESGs on the same LUN to allow future growth and expansion, thus creating a 520 GB LUN for DBs, and a 50 GB LUN for logs. For this building-block we will create a total of five DB LUNs and five log LUNs.
- ESGs: Will balance users between 10 storage groups

Step 3. Identify Exchange VM requirements

With the hardware available (server with four quad cores and 128 GB RAM) we can place 4,000 users per VM, and will need 24 GB RAM per each VM to support this user profile. The following formula was used to calculate the required memory:

$$4000 * 5 \text{ MB} + 5 \text{ GB} = 24 \text{ GB RAM}$$

Step 4. Finalize and define your Exchange VM building block

Building-block 1 was created by applying the criteria outlined in steps 1 through 3 above, as follows.

Building-block 1

No. of users per Exchange Mailbox Server/VM	User IOPS profile	Mailbox size (MB)	vCPUs per VM	Memory per VM (GB)	Disk configuration	ESGs per Exchange VM
4,000	0.48	500	4	24	Sixteen for databases and four for logs	10

Validation of the virtualized Exchange 2007 environment

The scope of our performance validation testing involved:

- Using JetStress to identify attainable IOPS for the CLARiiON CX4-960 (based on 64,000 Exchange users, with building-block 1 for configuring Exchange VMs)
- Using JetStress to validate CX4-960 storage array performance (based on a combination of different Exchange user profiles—configuration with 44,000 Exchange users)
- Simulating a failure of the CLARiiON SP during JetStress test
- Using LoadGen to validate both the Exchange environment e-mail flow and VMware configurations

Testing methodologies

To ensure that the building-block approach can scale up in a large Exchange deployment environment, we have performed end-to-end validation of the entire infrastructure by using standard recommended tools for Exchange pre-deployment testing.

Exchange performance validation tools

Performance validation tools used to test the environment included:

- **Microsoft Exchange JetStress**—Verifies the performance and stability of the disk subsystem before putting the Exchange server into a production environment.
- **Microsoft Load Generator (LoadGen)**—LoadGen is used to perform benchmarking, pre-deployment validation, and stress testing tasks that introduce various types of workloads into a test (non-production) Exchange messaging system. It simulates the delivery of multiple MAPI, Outlook Web Access, IMAP, POP, and SMTP client messaging requests to an Exchange server. LoadGen requires the Exchange environment to be deployed for validation testing (non-production configuration).

Performance collection and data analysis tools

Performance collection and data analysis tools used to test the environment included:

- **CLARiiON Navisphere Analyzer**—EMC Navisphere Management enables you to manage, discover, monitor, and configure your CLARiiON storage systems.
- **Windows Performance Monitor (perfmon)**—The Windows Performance Monitor (also known as system monitor) is a Performance Diagnostic snap-in for the Microsoft Management Console (MMC). From this single console you can track a range of system processes and obtain real-time graphical displays of test results. You can customize the data you want to collect in logs by defining thresholds for alerts and automatic actions, generate reports, and view past performance data.
- **VMware esxtop**—The esxtop command line tool provides a granular look at how ESX Server uses resources in real time. The tool runs on the ESX Server host's service console.

CLARiiON CX4-960 performance validation with JetStress based on the Microsoft ESRP testing framework

Microsoft JetStress was used to validate the CLARiiON CX4-960 storage subsystem and its capability to handle a very heavy Exchange I/Os load. The testing performed is similar to the one required for Microsoft's ESRP program that is designed for storage vendors like EMC to submit their storage solutions for Exchange 2007 deployments.

Microsoft ESRP testing overview

The Microsoft's ESRP program provides a common storage testing framework for vendors to provide information on its storage solutions for Microsoft Exchange Server software. The ESRP program is not designed to be a benchmarking program; tests are not designed to measure the maximum throughput for a given solution. ESRP testing focuses on producing recommendations from vendors for supported Exchange configurations. For more details on the Microsoft ESRP storage program see:

<http://technet.microsoft.com/en-us/exchange/bb412164.aspx>

Note: Currently, Microsoft does not accept ESRP submissions from storage vendors if Exchange Mailbox Servers are configured in a virtualized environment (VMware or Microsoft Hyper-V). That is why EMC finds it important to provide customers with vital data points and performance information to help deploy or transition an Exchange environment from physical to virtual when you are ready.

Test results listed in this white paper are provided according to ESRP guidelines for easy comprehension and comparison to other ESRPs published on Microsoft's web site (including ESRP submissions from EMC and other storage providers).

JetStress testing phase 1 summary

In this phase, testing was performed to identify the attainable load that the CLARiiON CX4-960 can sustain by handling maximum sustainable Exchange-type I/O for 2 and 24 hours. For this phase we have configured 16 Exchange VMs (load-balanced between four ESX Servers) to configure building-block 1. This resulted in a total configuration with 64,000 very heavy Exchange users.

Table 4 and Figure 2 detail the configuration of the environment for the 64,000-users simulation.

Simulated Exchange configuration for JetStress testing phase 1

Table 4 lists the simulated Exchange configuration details.

Table 4. Simulated Exchange configuration for phase 1

Item	Description
Number of Exchange mailboxes simulated	64,000
Number of physical hosts (VMware ESX Servers)	4
Number of virtual Exchange hosts (VMs)	16 (four per ESX Server)
Number of mailboxes per Exchange host (VM44,000)	4,000
Number of storage groups per Exchange hosts (VMs)	10
Number of mailbox stores/storage group	1
Number of mailboxes/mailbox store	400
Number of mailbox store LUNs/storage group	1
Simulated profile: I/Os per second per mailbox (IOPS include 20% headroom)	0.48
Database LUN size	520 GB (two databases per LUN)
Log LUN size	50 (two log volumes per LUN)
Backup LUN size/storage group	N/A
Total database size for performance testing	32,000
% storage capacity used by Exchange database	62.5%

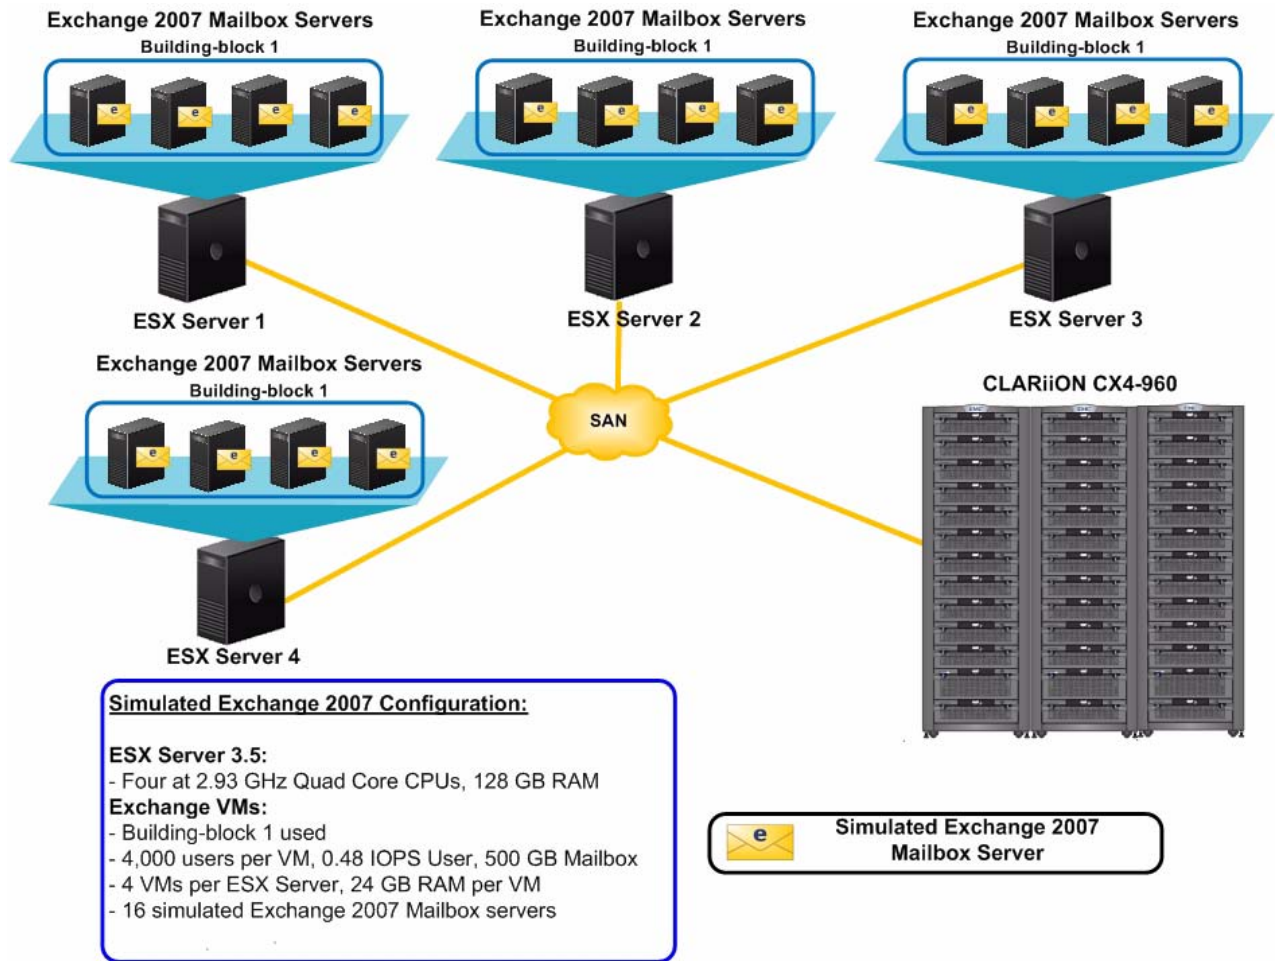


Figure 2. Exchange 2007 configuration for JetStress validation with 64,000 users

Primary storage hardware

Table 5 lists the primary storage hardware used in the environment.

Table 5. Primary storage hardware

Item	Description
Storage connectivity (Fibre Channel, SAS, SATA, iSCSI)	Fibre Channel
Storage model and OS/firmware revision	CLARiiON CX4-960, FLARE OS 4.28; WHQL listing: http://www.windowsservercatalog.com/item.aspx?idItem=37744d95-b6d7-235d-76d9-36f00352b9c2&bCatID=1282
Storage cache	32 GB
Number of storage controllers	2
Number of storage ports	8
Maximum bandwidth of storage connectivity to host	4 GB/s
Switch type/model/firmware revision	Cisco MDS 9509 / 4 GB/s, Fibre Channel switch
HBA model and firmware	QLogic (QLA2432)
Number of HBAs per ESX Server	4
Physical ESX host server type	Dell R900, 4 Quad-core CPUs @ 2.93 GHz, 128 GB RAM
Total number of disks tested in solution	320
Maximum number of spindles that can be hosted in the storage	960

Primary storage software

Table 6 lists the primary storage software used in the environment.

Table 6. Primary storage software

Item	Description
Multipathing	Native ESX 3.5 multipathing
VM OS	Windows 2003 R2 X64 Enterprise with SP2
ESE.dll file version	8.1.340.1

Primary storage disk configuration (mailbox store disks)

Table 7 lists the primary storage disk configuration (mailbox store disks) for the environment.

Table 7. Primary storage disk configuration (mailbox store)

Item	Description
Disk type, speed and firmware revision	FC, 15k rpm
Raw capacity per disk (GB)	402
Number of physical disks in test	256
Total raw storage capacity (GB)	102,912 GB
Number of disks per LUN	8
RAID level	RAID 1_0
Total formatted capacity	41,600 GB
Storage capacity utilization	40%
Database capacity utilization	31%

Primary storage disk configuration (transactional log disks)

Table 8 lists the primary storage disk configuration (transactional log disks) for the environment.

Table 8. Primary storage disk configuration (transactional log disks)

Item	Description
Disk type, speed and firmware revision	FC, 15k rpm
Raw capacity per disk (GB)	133
Number of spindles in test	64
Total raw storage capacity (GB)	8,512
Number of disks per LUN	4
Number of slices per LUN or number of disks per LUN	4
RAID level	RAID 1_0
Total formatted capacity	4,000 GB

Performance testing across server metrics

This section provides a high-level summary of the test data generated by the Microsoft ESRP testing framework.

Reliability

A number of the tests in the Microsoft ESRP framework are designed to test reliability over a 24-hour period. The goal of these tests is to verify that the storage can handle a high I/O load for a long period of time. Following the stress test, both log and database files are analyzed for integrity to ensure there is no database/log corruption. After analysis of the test results the following was observed:

- No errors were reported in the event log file for the storage reliability testing.
- No errors were reported for the database and log checksum.

Storage performance results

Performance testing exercises the storage with maximum sustainable Exchange-type I/O for two hours. The test is used to show how long it takes for the storage to respond to an I/O under load. [Table 9](#) provides a sample taken from each of the simulated VM hosts. It is the average of all the logical disks in the two-hour test duration. All 16 VMs in the test are listed together, and the aggregate numbers across all servers are listed in [Table 10](#). See "[CLARiiON CX4-960 performance](#)" for additional CLARiiON performance data.

Individual server metrics

Table 9 shows individual server JetStress performance results for all 16 virtualized servers in the simulation test. The test was based on the following parameters:

- User profile (IOPS): 0.48
- Mailbox size: 500 MB
- Users per host (per VM): 4,000

Table 9. Individual server metrics—phase 1

Server	Target IOPS	Database I/O				Transaction log I/O		
		Database disk transfers/sec (IOPS)	Average database disk write latency (ms)	Average database disk read latency (ms)	Average database disk reads/sec	Average database disk write/sec	Average log disk writes/sec	Average log disk write latency (ms)
MBX1	1920	2,043	0.016	0.010	110.137	94.173	30.706	0.003
MBX2	1920	2,068	0.016	0.009	111.265	95.537	31.203	0.003
MBX3	1920	2,072	0.016	0.009	111.561	95.623	31.191	0.003
MBX4	1920	2,103	0.016	0.009	113.285	96.996	31.493	0.003
MBX5	1920	1,999	0.016	0.010	107.700	92.217	30.576	0.003
MBX6	1920	2,564	0.013	0.008	138.036	118.385	38.727	0.002
MBX7	1920	1,998	0.016	0.010	107.792	91.991	30.519	0.003
MBX8	1920	2,432	0.013	0.008	131.141	112.048	37.016	0.002
MBX9	1920	2,060	0.016	0.010	110.784	95.219	31.411	0.003
MBX10	1920	2,078	0.016	0.010	111.914	95.889	31.494	0.003
MBX11	1920	2,657	0.012	0.007	143.424	122.275	40.335	0.002
MBX12	1920	1,970	0.016	0.010	106.155	90.864	30.144	0.003
MBX13	1920	2,008	0.016	0.010	108.252	92.566	30.720	0.003
MBX14	1920	1,991	0.016	0.010	107.412	91.659	30.364	0.003
MBX15	1920	2,021	0.016	0.010	109.102	93.018	30.928	0.003
MBX16	1920	2,043	0.016	0.009	109.895	94.431	31.260	0.003
Total IOPS	30,720	34,107						

Aggregate performance across all servers metrics

Table 10 lists the sum of I/Os across servers in this configuration and the average latency across all servers. The total IOPS achieved exceeded the 30,720 number of IOPS required to support this user configuration (64,000 users * 0.48 IOPS).

Table 10. Aggregate performance across all servers metrics

Database I/O	Value
Database disks transfers/sec (achieved IOPS)	34,107
Target Exchange IOPS	30,720
Database disks reads/sec	1837.8
Database disks writes/sec	1572.9
Average database disk read latency (ms)	15
Average database disk write latency (ms)	9

Transaction log I/O	Value
Log disks writes/sec	518
Average log disk write latency (ms)	3

JetStress testing phase 2 summary

In this phase, testing was performed to identify CLARiiON CX4-960 performance by executing an Exchange user load with a combination of multiple building-blocks and user profiles (with 500 MB and 2 GB mailboxes and user profiles at 2 IOPS). This configuration included a total of 12 VMs with a total user configuration of 44,000 users. Building-blocks 1, 2 and 3 were used for this simulation.

The environment for JetStress phase 2 testing was configured as follows:

- 8 VMs with building-block 1 (32,000 users)
- 2 VMs with building-block 2 (8,000 users)
- 2 VMs with building-block 3 (4,000 users)

Table 11 details how the simulated Exchange environment was configured for 44,000 users. In this configuration, RAID 5 was also tested for large mailboxes (building-block 2) in order to provide the point of reference and performance data.

Simulated Exchange configuration for JetStress testing phase 2

Table 11 lists the simulated Exchange configuration details.

Table 11. Simulated Exchange configuration for phase 2

Item	Description
Number of Exchange mailboxes simulated	44,000
Number of physical hosts (VMware ESX Servers)	3
Number of virtual Exchange hosts (VMs)	12
Number of mailboxes per Exchange hosts (VMs)	4,000 per VM for BB1 (8 VMs - 32,000 users) 4,000 per VM for BB2 (2 VMs - 8,000 users) 2,000 per VM for BB3 (2 VMs - 8,000 users)
Number of storage groups per Exchange host (VM)	10 ESGs for BB1 40 ESGs for BB2 (VM 9-10) 20 ESGs for BB3 (VM 11-12)
Number of mailbox stores/storage group	1
Number of mailboxes/mailbox store	400 users for BB1 100 users for BB2 100 users for BB3
Number of mailbox store LUNs/storage group	1
Simulated profile: I/Os per second per mailbox (IOPS include 20% headroom)	0.48 IOPS for BB1 and BB2 2 IOPS for BB3
Database LUN size	520 GB (Two databases per LUN)
Log LUN size	50 GB for BB1 (Two log volumes per LUN) 30 GB for BB2 (Two log volumes per LUN) 40 GB for BB3 (Two log volumes per LUN)
Total database size for performance testing	40,000 GB
% storage capacity used by Exchange database	76.9%

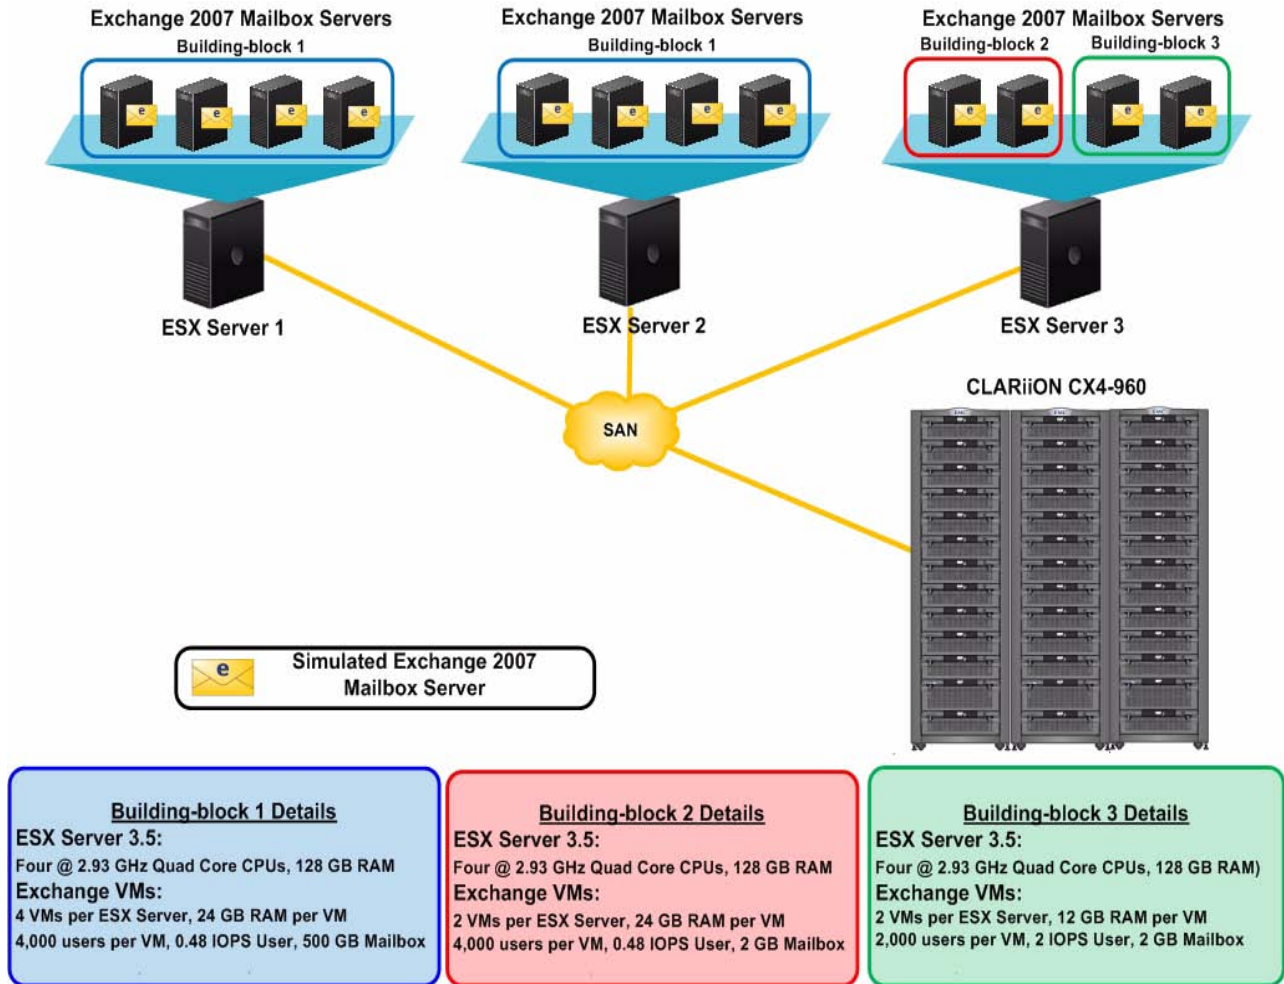


Figure 3. Exchange 2007 configuration for JetStress validation with 44,000 users

CLARiiON storage hardware used in this configuration

Table 12 lists the CLARiiON total number of disks used in the JetStress phase 2 test environment. For a complete list of CLARiiON storage hardware used; see Table 7. The only difference between the two phases is the number of spindles.

Table 12. CLARiiON storage hardware used in this configuration

Item	Description
Total number of disks tested in configuration	310

Disk configurations

Table 13 and Table 14 list the disk configurations for mailbox store and transaction log disks for the environment.

Table 13. Disk configuration (mailbox store disks)

Item	Description
Disk type, speed and firmware revision	FC, 15k rpm
Raw capacity per disk (GB)	402
Number of physical disks in test	262
Total raw storage capacity (GB)	105,324
Number of disks per LUN	8 for BB1 and BB3, 5 for BB2
RAID level	RAID 1_0 for BB1 and BB3, RAID 5 for BB2
Total formatted capacity	52,000 GB
Storage capacity utilization	49.3%
Database capacity utilization	37.9%

Table 14. Disk configuration (transactional log disks)

Item	Description
Disk type, speed and firmware revision	FC, 15k rpm
Raw capacity per disk (GB)	133 for BB1, 402 for BB2 and BB3
Number of spindles in test	48
Total raw storage capacity (GB)	10,710
Disk slice size (GB)	N/A

Table 14. Disk configuration (transactional log disks) (continued)

Item	Description
Number of disks per LUN	4
RAID level	RAID 1_0
Total formatted capacity	4,000 GB

Individual server metrics

Information in these tables includes the sum of I/Os across storage groups and the average latency across all storage groups on a per-server basis. [Table 15](#) shows individual server performance JetStress results for all 12 virtualized servers in the simulation test. It is important to point out that the performance achieved is well within Microsoft's guidelines for Exchange 2007 subsystem performance.

Table 15. Individual server metrics—phase 2

Server	User profile (IOPS)	Mailbox size (MB)	Users per host (per VM)	Target IOPS	Database I/O			Transaction log I/O			
					Database disk transfers/sec (IOPS)	Average database disk write latency (ms)	Average database disk read latency (ms)	Average database disk reads/sec	Average database disk write/sec	Average log disk writes/sec	Average log disk write latency (ms)
MBX1	0.48	500	4,000	1,920	2,220	0.015	0.009	243.732	200.301	129.765	0.002
MBX2	0.48	500	4,000	1,920	2,231	0.014	0.009	244.902	201.295	130.112	0.002
MBX3	0.48	500	4,000	1,920	2,231	0.015	0.009	244.832	201.397	130.126	0.002
MBX4	0.48	500	4,000	1,920	2,292	0.014	0.009	251.215	207.260	132.832	0.002
MBX5	0.48	500	4,000	1,920	2,829	0.012	0.007	309,737	256.006	164.301	0.002
MBX6	0.48	500	4,000	1,920	2,269	0.014	0.009	248.765	205.011	131.535	0.002
MBX7	0.48	500	4,000	1,920	2,285	0.014	0.009	250.286	206.793	132.478	0.002
MBX8	0.48	500	4,000	1,920	2,787	0.012	0.008	305.552	251.928	162.014	0.002
MBX9	0.48	2,000	4,000	1,920	2,575	0.015	0.009	73.998	54.732	42.790	0.002
MBX10	0.48	2,000	4,000	1,920	2,564	0.014	0.009	73.639	54.565	42.569	0.002
MBX11	2.0	2,000	2,000	4,000	4,256	0.015	0.011	230.384	195.244	127.066	0.002
MBX12	2.0	2,000	2,000	4,000	4,267	0.015	0.011	230.937	195.762	127.063	0.002
Total IOPS					27,200	3,2807					

Aggregate performance across all server metrics

The sum of I/Os across servers in this configuration and the average latency across all servers are provided in Table 16. The total of 32,807 IOPS achieved exceeded the 27,200 number of IOPS required to support this configuration. Figure 4 shows individual server performance and IOPS achieved during the 24-hour stress test.

Table 16. Aggregate performance across all servers metrics

Database I/O	Value
Database disks transfers/sec (achieved IOPS)	32,807
Target Exchange IOPS	27,200
Database disks reads/sec	2,707.9
Database disks writes/sec	2,230.3
Average database disk read latency (ms)	14
Average database disk write latency (ms)	9

Transaction log I/O	Value
Log disks writes/sec	1,452.6
Average log disk write latency (ms)	2

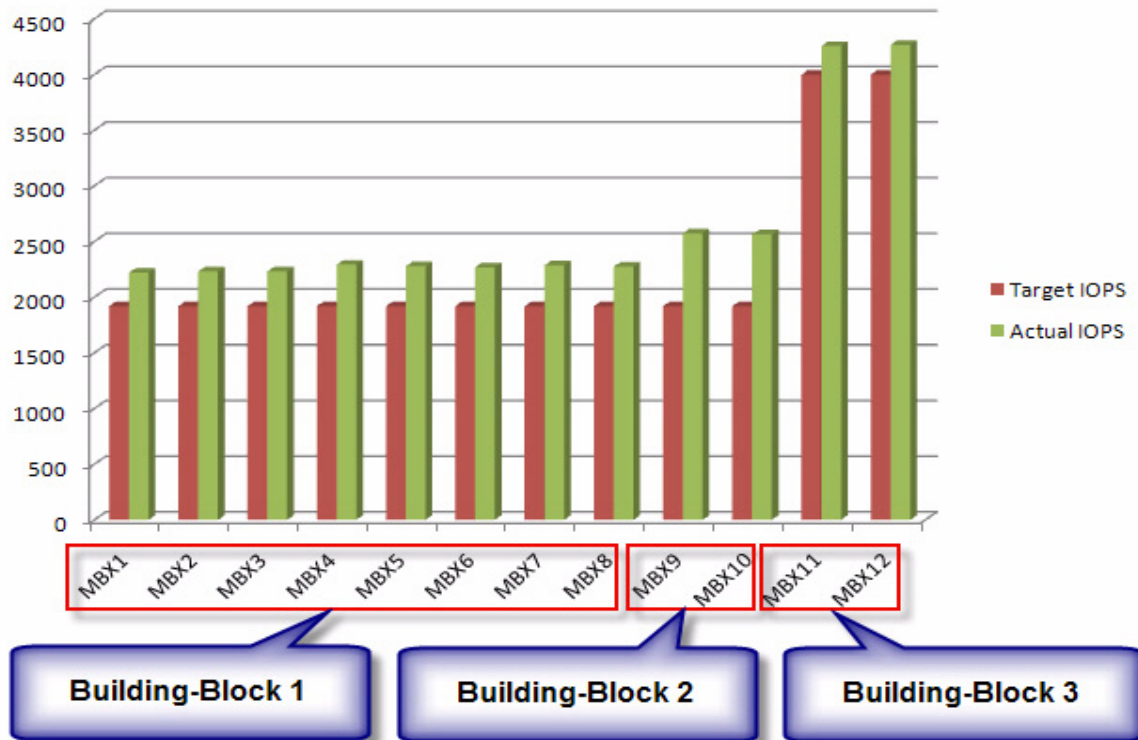


Figure 4. 24-hour JetStress test results—Exchange IOPS achieved per building-block 1, 2, 3

Streaming backup/recovery performance

There are two tests in this section. These tests simulate Exchange streaming backup, and measures the streaming read performance of the database disks. These tests measure database and logs read I/O performance metrics by running checksum on all of the databases and log files.

Database read-only performance test

This test is used to measure the maximum rate at which databases could be recovered. [Table 17](#) shows the average rate for a single database file.

Table 17. Database Read-only performance test results

	Building-block 1 (MBX1 through MBX8)	Building-block 2 (MBX9 and MBX10)	Building-block 3 (MBX11 and MBX12)
MB read/sec per storage group	10.29	13.44	13.44
MB read/sec total	822.81	412.25	412.25

Log read-only performance test

This test is used to measure the maximum rate at which the log files can be played against the databases. [Table 17](#) shows the average rate for 500 log files played in a single storage group. Each log file is 1 MB in size.

Table 18. Log Read-only performance test results

	Building-block 1 (MBX1—MBX8)	Building-block 2 (MBX9 and MBX10)	Building-block 3 (MBX11 and MBX12)
Average time to play one log file (sec)	1.82	3.04	2.48

CLARiiON CX4-960 performance

During both phases of JetStress testing, CLARiiON performance was monitored and analyzed using Navisphere Analyzer. Figure 5 shows total I/O throughput achieved. Figure 6 shows total SP utilization under a simulated very heavy Exchange load.

In summary, the data shows that CX4-960 can comfortably handle various Exchange user loads and produced ~42,000 I/Os with SP utilization at approximately 80 percent.

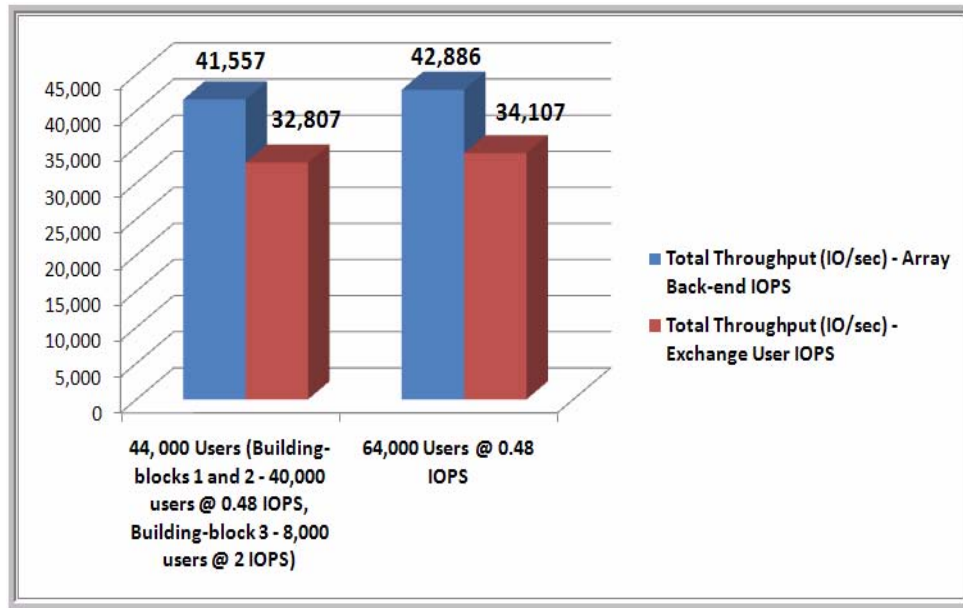


Figure 5. CLARiiON CX4-960 total throughput (IOPS) during JetStress testing

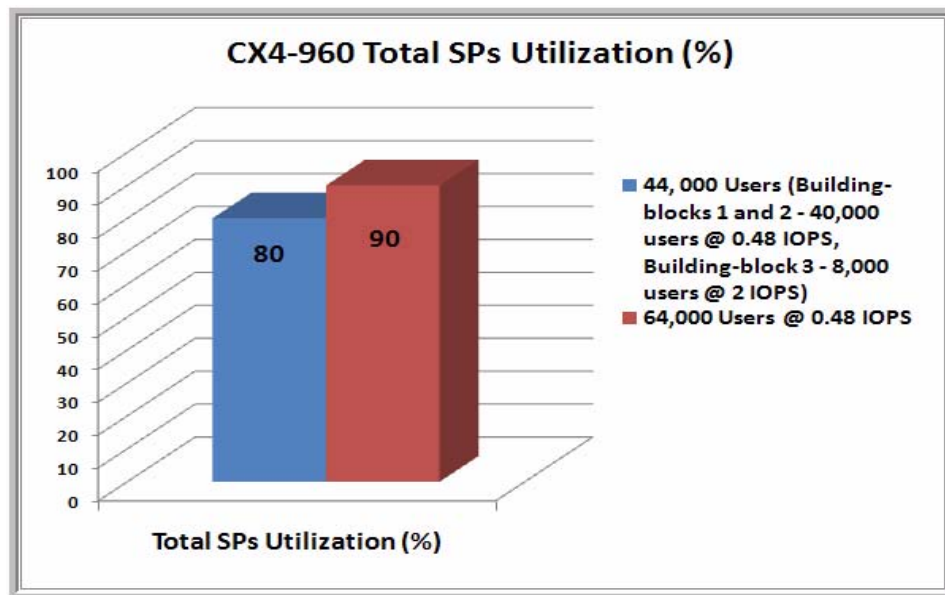


Figure 6. CLARiiON CX4-960 SP utilization during JetStress testing

JetStress testing – phase 3 summary

Many of our customers ask, what will happen if they lose an SP in their CLARiiON? In order to provide the answer and to alleviate customer fears, we have performed a test where an SP was removed from the CLARiiON array while JetStress was executing I/Os against the storage system.

JetStress testing Phase 3–CLARiiON storage processor simulated failure test

An SP pull test was performed while the 8-hour JetStress test was in progress to prove that Exchange 2007 users can sustain uninterrupted service from the CLARiiON CX4-960 in the unlikely event of an SP failure. During the test, **SP B** was removed from the array for 2 hours; see [Figure 7](#).

Test results showed no interruption in the I/O traffic between the CX4-960 and ESX Servers hosting Exchange VMs. The LUNs that were homed on the failed SP (**SP B**) automatically moved (trespassed) to another SP (**SP A**) and the CX4-960 continued to service client requests.

When **SP B** was re-inserted, the LUNs owned by that SP trespassed back automatically and I/O traffic was once again evenly balanced between the two SPs.

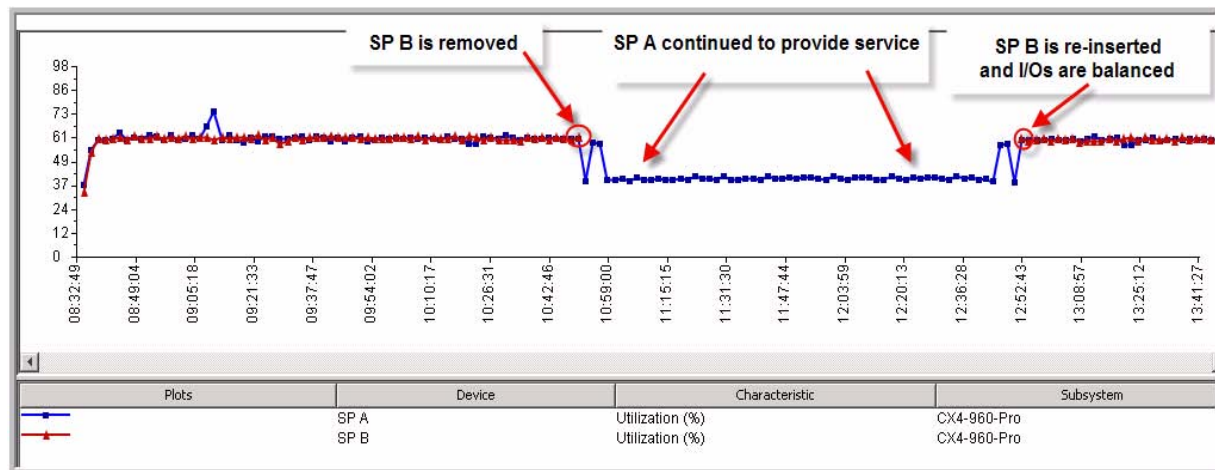


Figure 7. CLARiiON SP failure simulation test

Exchange 2007 environment validation with LoadGen

As was described earlier in this paper, the Microsoft LoadGen tool simulates the impact of numerous MAPI clients on the Exchange servers. To perform this testing, the Exchange 2007 environment was fully deployed with all of its components and user mailboxes were created as per each Exchange VM building-block user profile we have defined earlier.

Why is LoadGen testing necessary?

As opposed to JetStress (which only validates underlying storage) Microsoft's LoadGen performs end-to-end validation of all aspects of the entire Exchange infrastructure and virtualization layer components such as:

- Storage array
- SAN/Ethernet
- VMware ESX Server
- Exchange Mailbox VMs
- Exchange Hub Transport and Client Access servers (HUB/CAS)
- Active Directory Domain Controllers/Global Catalog servers

Measuring good Exchange performance

Good Exchange performance can be measured by multiple factors such as overall server utilization, disk latencies and good response times from Outlook clients (RPC latencies). Some of the performance parameters analyzed during LoadGen testing included:

- **RPC requests**

The RPC Requests performance counter indicates the number of MAPI remote procedure calls (RPC), or client requests that are currently being processed by the Microsoft Exchange Information Store service (tasks such as sending e-mail, browsing the calendar, and others). The average value of the RPC Requests performance counter should be under 30 at all times.

- **RPC average latency**

The RPC Average latency counter identifies the time the server takes to process RPC Requests. This value should always be lower than 50 milliseconds (for organizations with mostly online mode clients) or 100 milliseconds (for organizations with mostly cached mode clients).

- **Disk latencies**

Disk latencies are very important for the Exchange application and a healthy underlying storage. Follow Microsoft recommendations for disk latencies in the following article:

<http://technet.microsoft.com/en-us/library/bb124518.aspx>

- Disk latencies for our testing included PhysicalDisk(Instance)Avg. Disk sec/Read and PhysicalDisk(Instance)/Avg. Disk sec/Write; see [Table 15](#) for individual server performance based on these counters.

Disk latencies indicate the average time (in milliseconds) to read data and write from the disk. The average value should be below 20 ms with spikes (maximum values) no higher than 50 ms, that last for more than a few seconds.

- **Exchange Mailbox Servers, CPU utilization**

As seen from numerous tests performed, for a steady performance Exchange server VM, CPU utilization should stay below 50-60 percent when under steady heavy load. If a higher CPU utilization is noticed over a long period of time, verify total IOPS generated (total disk transfers performance counter) and make adjustments as necessary.

- **ESX Server overall performance**

Monitoring ESX Server performance is a key to a healthy Exchange environment. Follow VMware recommendation and best practices for monitoring the ESX Server and pay attention to the following parameters to ensure that there are no bottlenecks:

- Overall CPU utilization
- Memory usage
- Networking utilization

LoadGen testing—environment configuration

The Exchange environment for LoadGen testing was configured based on the same layout used in JetStress phase 2 testing (44,000 users with multiple building-blocks configured on 12 Exchange Mailbox Server VMs). In addition, another ESX Server was deployed to host all the peripheral components. **Figure 8** depicts a fully deployed Exchange environment with all of its components, including HUB/CAS servers, Domain Controller servers, and an EMC Replication Manager backup server.

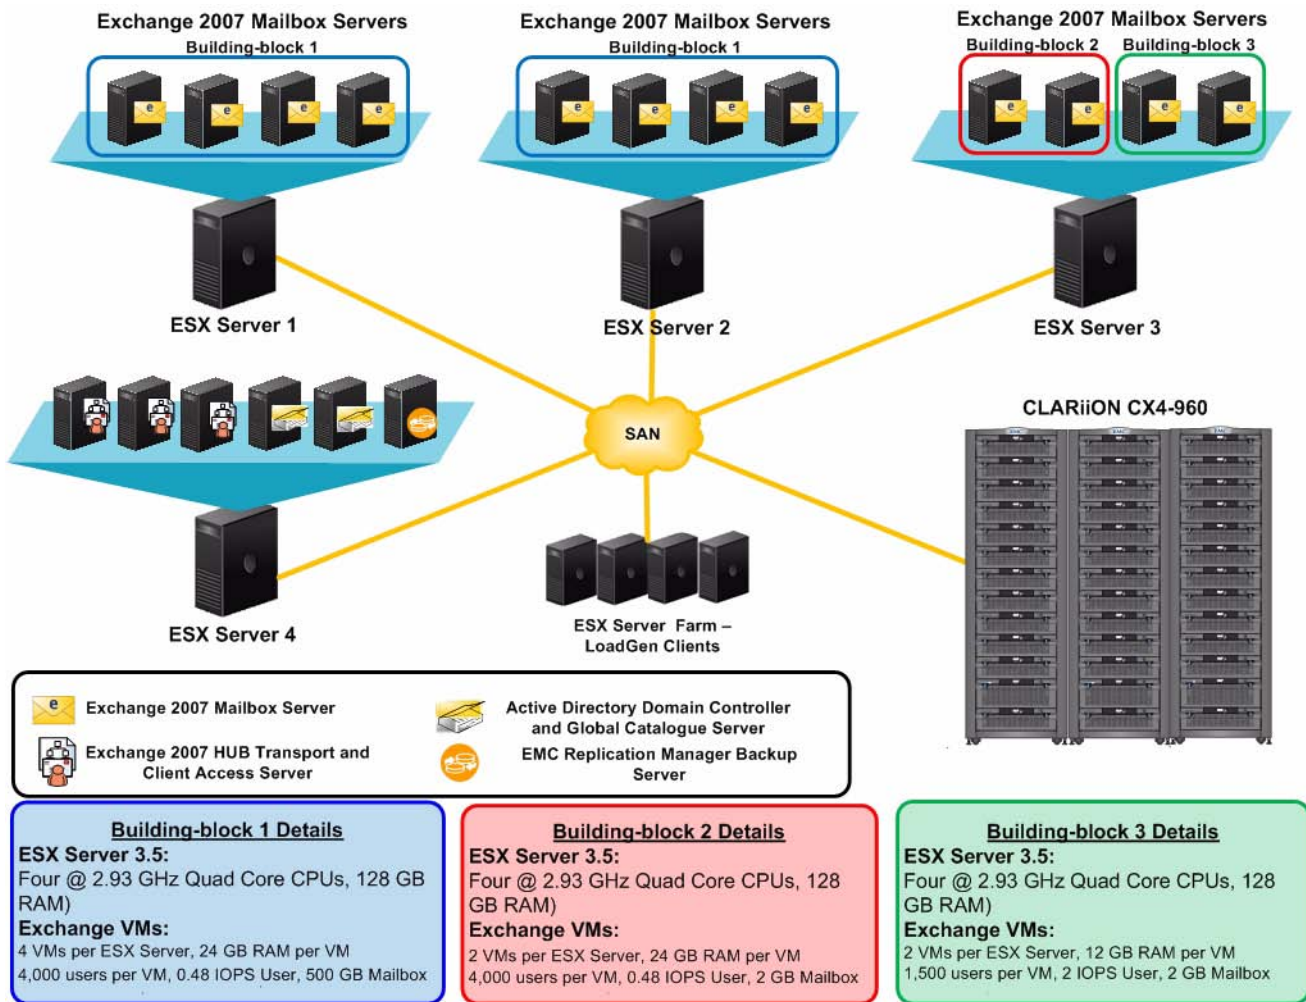


Figure 8. Exchange environment configuration for LoadGen validation testing (44,000 users)

During each LoadGen simulation test, Exchange 2007 tasks were dispatched to Exchange users according to the parameters specified in the LoadGen workload configuration file; see [Table 19](#). Based on the defined user profile, the load was executed over the duration of the test, thus generating Exchange 2007 logs. After each simulation test completed, a final HTML summary report was generated that details the total number of tasks dispatched and tasks completed. All of this information along with the performance monitor counters was used to analyze the test results.

Table 19. LoadGen workload profile configuration parameter details

LoadGen parameter	Value
Simulated day	Varies from 4 to 8 hours
Total length of simulation	Varies from 4 to 8 hours
Stress mode	Disabled
Distribution lists	100
Dynamic distribution lists	None
Contacts	None
Action profile	Very Heavy (BB1 and BB2) Custom (BB3)
User mode	Outlook 2007 online
Messages per inbox	BB1—10,000 BB2—4,500 (with custom messages) BB3—4,500 (with custom messages)
Custom messages used	For BB2 and BB3 (messages from 50 KB to 5 MB in size with attachments)

Virtualized Exchange environmental performance

Performance results achieved during LoadGen validation are provided in [Table 20](#). [Figure 9](#) and [Figure 10](#) show individual Exchange VM performance results during simulation.

As shown, great overall performance of the entire virtualized Exchange environment was achieved. This is measured by each of the building-block's performance characteristics in relation to the load (total IOPS achieved) that each Exchange VM was intended to carry (and the load that was produced by the LoadGen simulation).

For example, Exchange VMs configured with 4,000 users (building-block 1 with user profile @ 0.48 IOPS) exhibited steady utilization of 48 percent while handling 2,000 IOPS. At the same time, the overall ESX Server utilization handling all four VMs (16,000 users) configured with the same user profile performed at 75 percent; see [Figure 11](#).

Table 20. Exchange and VMware ESX performance with LoadGen

Performance parameter	Building-block 1 4,000 users 0.48 IOPS 500 MB mailbox	Building-block 2 4,000 users 0.48 IOPS 2 GB mailbox	Building-block 3 2 IOPS, 2 GB mailbox	
			1,500 users	2,000 users
RPC averaged latency	4	6	5	7
RPC requests	2	7	9	12
Total disk transfers achieved during simulation (IOPS)	~2,000	~2,200	~3,100	~4,100
Outlook tasks completed per server (send/receive/browse calendar, post free busy, etc.)	715,084	715,015	1,370,664	1,826,028
Transaction logs generated during simulation (per Exchange VM)	~48,000	~48,000	~45,000	~60,000
Transaction logs generated during simulation (per user)	12	12	30	35
VMs CPU utilization (%)	48%	55%	40-45%	65-70%
ESX Server CPU utilization	75%	83%	83%	83%

As seen in [Table 20](#), building-block 3 performance is characterized by different user loads (1,500 users and 2,000 users, respectively). Two thousand users at 2 IOPS using four vCPU cores (500 users per core) produced more than 4,000 disk transfers/sec (IOPS) and pushed Exchange VM utilization to 70 percent. Although this server utilization can be acceptable by some customers, others would prefer to see a lower utilization (less than 50 percent) that can be achieved by reducing the number of users per VM from 2,000 to 1,500. In that case, less IOs (~3100) and less users per core (375) will drive Exchange VM CPU utilization to 45 percent.

Figure 9 shows an Exchange VM under load configured using the building-block 3 user profile of 2 IOPS. It can be clearly seen how a reduction in users affected VM CPU utilization. That is why we recommend that you follow users-to-IOPS-to-CPU cores guidelines. Refer to "Exchange VM questions for planning" when sizing your Exchange VMs.

Note: With newer hardware and faster CPUs it may be possible to place more users on each VM. With VMware ESX 4.0 it is possible to allocate up to eight CPUs per VM, thus increasing the amount of Exchange users per VM.

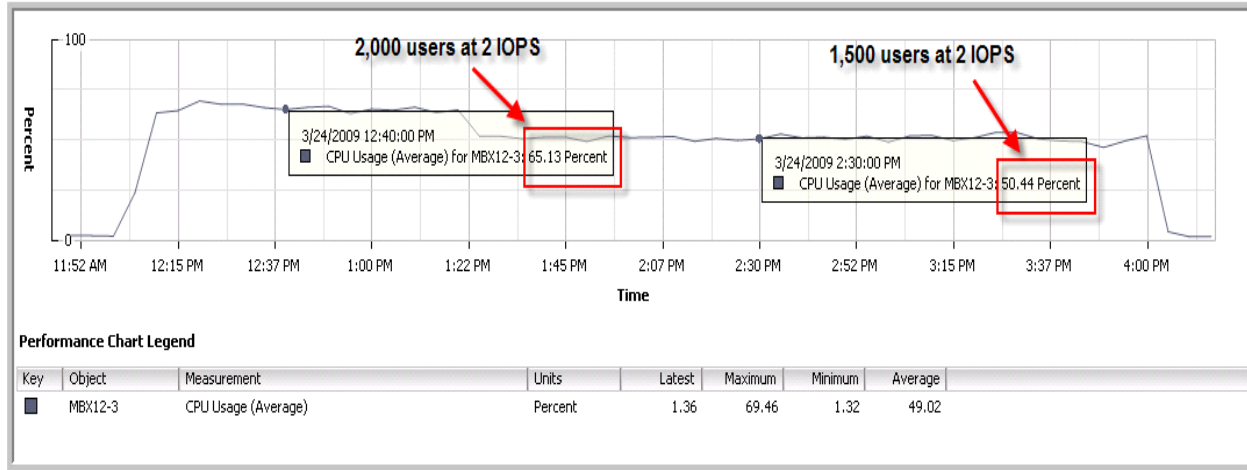


Figure 9. Exchange Server VM utilization with 2,000 and 1,500 users at a 2 IOPS user profile

Figure 10 and Figure 11 show individual Exchange VM performance results during simulation. More specifically, Figure 10 shows CPU utilization of Exchange server VMs configured with user profiles defined in building-block 1 (500 MB mailbox, 0.48 IOPS). A very stable, even utilization (around 48 percent) was noticed across all four VMs on a single ESX Server.




Figure 10. Building-block 1 validation—Exchange VMs performance

Figure 11 shows CPU utilization of the Exchange server VMs configured with user profiles defined in building-block 2 (2 GB mailbox, 0.48 IOPS - MBX9 and MBX10) and building-block 3 (2 GB mailbox, 2 IOPS - MBX11 and MBX12). The test shows that the VMs were capable of handling large spikes at the beginning of the test while memory was still not optimized and stabilized (45-56 percent) for the duration of the LoadGen simulation. Test results proved that VMs were sized correctly for the load they were intended to carry.




Figure 11. Building-block 2 and 3 validation—Exchange VMs performance

Exchange HUB/CAS performance

During LoadGen simulation, we observed very stable performance of virtualized HUB/CAS servers. Average CPU usage was no more than 20 percent for each server, and each HUB server handled about 16 messages per second while keeping a reasonable delivery queue. Figure 12 shows how many e-mails were processed and delivered via HUB Transport servers. It shows steady delivery without any bottlenecks at about 1,100 items per second for each server. Overall, all three HUB/CAS servers processed approximately 1.2 million messages during the 8-hour LoadGen simulation; see Figure 12.




Figure 12. Exchange HUB/CAS server performance

All performance data gathered and analyzed in this white paper, will help you size your Exchange VMs. Some of the important performance points gathered during the testing process, such as heavier I/O loads with a 2 IOPS profile, bring to light very important aspects of proper validating of the virtualized Exchange design.

EMC recommends to always validate your configuration with both tools if possible—JetStress and LoadGen. Based on your test results, necessary adjustments should be made before deploying Exchange in production.

Conclusion

We hope the information provided in this white paper will aid you when you are ready to virtualize your Exchange environment. The building-block approach and methodologies used to consolidate Enterprise Exchange 2007 on the EMC CLARiiON CX4-960 storage array and VMware ESX 3.5 will help you to deploy Exchange more easily and efficiently.

Rigorous performance and validation testing performed and discussed in this white paper demonstrates how virtualization technology can provide multiple benefits and provide the same level of performance with better utilization of the physical hardware.

EMC can help accelerate assessment, design, implementation, and management while lowering the deployment risks and cost of creating a virtualized Microsoft Exchange Server 2007 infrastructure. As mentioned earlier in the paper, to aid in virtualization decisions there is a new EMC service called Readiness Assessment for Virtualized Microsoft Applications; see "[Readiness Assessment for Virtualized Microsoft Applications](#)." This service involves conducting a thorough evaluation of your existing environment and requirements, and delivers detailed recommendations on how to optimize the deployment of your new virtualized configuration.

To learn more about this, and other validated solutions, visit the Resource Library on <http://www.emc.com>.

References

EMC has produced a growing library of technical documentation that details the testing and specific best practice guidelines for its virtualized Exchange solutions on CLARiiON storage. All of these documents are available on Powerlink, the EMC extranet community, for customers and partners.

The following EMC documents provide relevant information. Access to these documents is based on your login credentials to <http://powerlink.emc.com>:

- *Secure and Consolidated 16,000 Exchange Users Solution on a VMware/EMC Environment—Applied Technology* white paper
- *EMC CLARiiON Storage Solutions: Exchange 2007 Best Practices Planning* white paper

In addition, the VMware web page for Exchange virtualization includes the following technical documentation:

- *SAN System Design and Deployment Guide*
- *Performance Tuning Best Practices for ESX Server 3*
- *CLARiiON Integration with VMware ESX Server*
- *VMware Certified Compatibility Guides*

Also refer to the following Microsoft Support Policies article:

- *Support for Microsoft Software in VMware Virtual Machines*