

Performance Validation and Test Results for Microsoft Exchange Server 2010 Enabled by EMC CLARiiON CX4-960

A Detailed Review

EMC Global Solutions

Abstract

The purpose of this white paper is to profile the performance of the EMC[®] CLARiiON[®] CX4-960 with Microsoft Exchange Server 2010. The white paper provides best practices and guidelines for simplifying an Exchange storage configuration using a building-block approach, presenting validated performance results for 80,000 very heavy Exchange 2010 users with a 1 GB mailbox quota.

April 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 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: H6918.1

Performance Validation and Test Results for Microsoft Exchange Server 2010
Enabled by EMC CLARiiON CX4-960
A Detailed Review

Table of Contents

Executive summary	4
Introduction.....	6
Technology.....	8
Environment	10
Exchange Server 2010 storage design.....	14
Exchange 2010 database maintenance.....	19
Best practices planning	23
Testing methodology and results	27
Conclusion.....	33

Executive summary

Overview

When Exchange administrators are faced with the task of implementing a plan for a new Exchange 2010 deployment, they need to have a solid understanding of what is necessary to achieve optimal performance, while understanding the true value that the new, enhanced version will bring to their organization.

EMC performed comprehensive testing in order to guide Exchange administrators and provide methodologies, guidelines, and best practices to deploy Exchange 2010 disk layouts on the EMC® CLARiiON® CX4-960 array. In the tested configuration, one simple building-block was created to validate the design methodology and Exchange input/output per second (IOPS) on the CX4-960. The building-block is based on a single, very heavy Exchange user profile, and a mailbox size of 1 GB.

The building-block represents an amount of resources that are required to support a specified number of users on a single Exchange 2010 Mailbox server role. To ensure that the Exchange performance is within recommended Microsoft guidelines, rigorous validation and testing were performed using recommended Exchange performance validation tools such as Microsoft Jetstress.

This white paper presents administrators with validated test results, and includes:

- Best practices for configuring the EMC CLARiiON CX4-960 for Exchange 2010.
- CLARiiON storage performance validation for maximum Exchange 2010 IOPS.
- Guidelines for simplifying Exchange 2010 storage design using a building-block approach based on an Exchange 2010 user profile.
- Analysis of Exchange 2010 background database maintenance (BDM).

Business case

By choosing to optimize the Exchange 2010 messaging system with intelligent storage solutions from EMC, a company can:

- Reduce e-mail server and storage costs
- Improve e-mail server performance and employee productivity
- Reduce complexity and maximize agility

To address these challenges, powerful midrange CLARiiON CX4 arrays can be deployed to provide the necessary performance, maximum data capacity, and multiple level data protection for Exchange Server 2010.

This white paper focuses on the reliable performance of the EMC CLARiiON CX4-960 array with Microsoft Exchange 2010 by consolidating a large amount of users in a single storage framework.

Performance validation overview

The performance validation lays the groundwork for any enterprise business planning to deploy Exchange 2010.

This white paper presents methodologies and best practices, based on validated test results for designing a Microsoft Exchange Server 2010 environment utilizing the CLARiiON CX4-960 storage array. The unique combination of a breakthrough hardware design and advanced software capabilities enables the CLARiiON CX4-960 system to meet the growing IT challenges of today's 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

For example, the simulated test environment implements a very reliable SAN infrastructure and physical mailbox configuration. It consists of eight simulated Exchange 2010 Mailbox servers with the maximum Microsoft-recommended user configuration limit of 10,000 users per server. This configuration is designed to be able to support two database copies on a single CLARiiON CX4-960 array.

In addition to advancements in storage, Exchange Server 2010 introduces significant improvements in its database allowing for the reduction of input/output (I/Os). This, in turn, permits the effective use of cheaper, less expensive drive types such as Serial Attached SCSI (SAS), and Serial Advanced Technology Attachment (SATA) for Exchange 2010 mailbox storage. At the same time, these improvements allow for more effective use of high-performance Fibre Channel (FC) drives in optimal configurations that provide a higher level of performance and storage efficiency. New disks with larger capacities can now be used in optimal RAID 5 configurations, thus allowing for deployment of larger mailboxes for Exchange users with fewer disks.

Key results

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

- The building-block approach used to design Exchange 2010 storage on CLARiiON CX4-960 (with RAID 5) satisfied all recommended performance guidelines as provided by Microsoft.
- The Exchange 2010 BDM analysis was performed, and guidelines are provided to ensure that the same level of performance is achieved when BDM is enabled.

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 CLARiiON CX4-960 storage with Exchange Server 2010.

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

- Design guidelines used to build this environment
- Performance test results
- CLARiiON CX4-960 storage configuration testing using Jetstress

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 implementing Microsoft Exchange Server 2010 solutions with EMC CLARiiON CX4-960 storage

Terminology

This white paper includes the following terms.

Term	Definition
BDM	Process of the Exchange database maintenance that involves checksumming both active and passive database copies.
CLARiiON storage processors (SP)	Each CLARiiON CX4-960 contains two 2.33 GHz Quad-Core Intel Xeon processors. This dual-storage processor architecture provides the power for enhanced performance, stability, and reliability. Each SP contains 16 GB of ECC-protected system memory.
Database Availability Group (DAG)	A DAG is 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.
EMC PowerPath®	EMC PowerPath is host-resident software that works with both CLARiiON and EMC Symmetrix® storage systems to deliver intelligent I/O path management. Using PowerPath, administrators can improve the server's ability to manage heavy storage loads through continuous and intelligent I/O balancing. PowerPath automatically configures multiple paths, and dynamically tunes performance as the workload changes. PowerPath also adds to the HA capabilities of the CLARiiON and Symmetrix storage systems by automatically detecting and recovering from server-to-storage path failures.
Microsoft Jetstress	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 specific number of users.

Technology

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

- Exchange Server 2010
 - CLARiiON CX4-960
-

Exchange Server 2010

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.

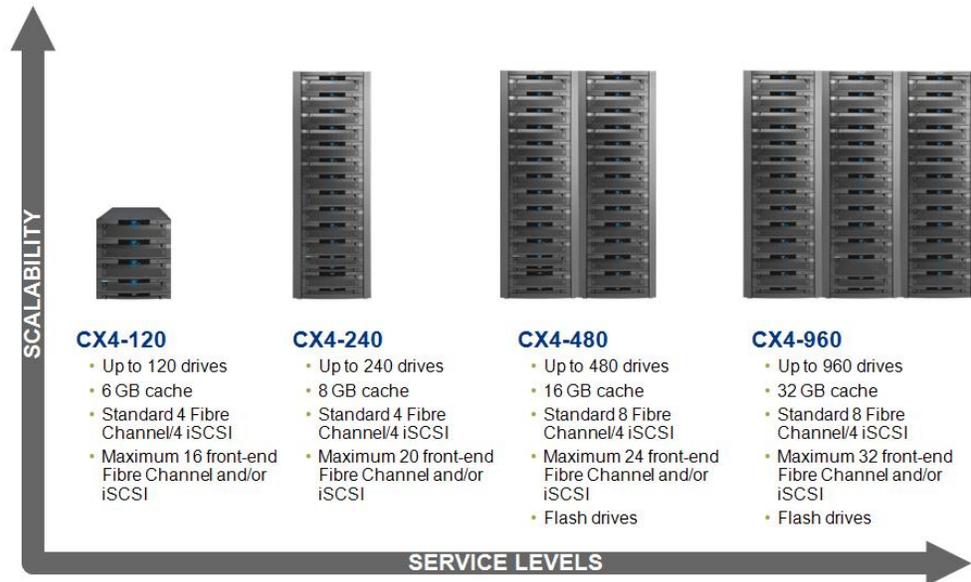
With the new version of Exchange 2010 Microsoft presents a new, unified approach to HA and disaster recovery (DR) by introducing features such as DAG, and online mailbox moves. Mailbox servers can now be implemented in mailbox resiliency configurations with database-level replication and failover.

Major improvements with the application database structure and I/O reduction include support for a larger variety of disk and RAID configurations including high-performance Enterprise Flash Drives (EFDs), FC drives, and slower-performing SATA and SAS drives.

**CLARiiON
CX4-960**

The EMC CLARiiON CX4-960 delivers performance, scalability, and advanced data management features in one, easy-to-use storage solution. It can support up to 512 highly available, dual-connected hosts and scales up to 960 disk drives for a maximum capacity of 1,899 TB.

The image below provides additional details about the EMC CLARiiON CX4 family.



The CLARiiON CX4-960 storage array integrates industry-changing midrange storage innovations that include:

- EFDs offering storage with 30 times more IOPS.
- UltraFlex™ technology for dual protocol connectivity, online expansion via I/O modules, and readiness for future technologies—such as 8 Gb/s 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—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.

Environment

Targeted customer profile

The simulated targeted customer profile is detailed in the following table.

Note

80,000 Exchange 2010 users is not the limiting factor for the CLARiiON CX4-960. The maximum user configuration will depend on the Exchange 2010 user profile and mailbox size.

For more information on determining your specific calculation, contact an EMC sales representative.

Item	Value
Number of users	80,000 users
Number of Exchange Mailbox servers	8 (physical)
Number of Exchange users per server	10,000
User profile (in Mailbox Resiliency configuration)	150 messages/user/day—0.15 IOPS
Mailbox size	1 GB
Database read/write ratio (in Mailbox Resiliency configuration)	3:2
Database maintenance configuration (24 x 7 BDM)	No BDM
Database copies (Mailbox Resiliency)	Two copies of each mailbox database

**Simulated
Exchange 2010
configuration**

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

Item	Value
Number of Exchange mailboxes simulated	80,000 mailboxes
Number of DAGs	1
Number of servers/DAG	8
Number of active mailboxes/server	10,000
Number of databases/hosts	16
Number of database copies	2
Number of mailboxes/database	625
Simulated profile: I/Os per second per mailbox (IOPS, includes 20% headroom)	0.15 (0.18 tested)
Database LUN size	820 GB
Log LUN size	65 GB
Total database size for performance testing	80,000 GB
Percent storage capacity used by the Exchange database	76.22%

Storage hardware

The hardware used in the environment is listed below.

Item	Description
EMC CLARiiON CX4-960	CLARiiON CX4-960, FLARE 29
Storage connectivity (FC, SAS, SATA, iSCSI)	FC
Storage cache	32 GB (600 MB Read cache, and 10,160 MB Write cache per SP)
Number of storage controllers	2
Number of storage ports	16 (8 per SP)
Maximum bandwidth of storage connectivity to host	16 * 4 Gb/s
FC switch	4 Gb/s enterprise-class FC switch (requires a minimum of 16 ports)
Host bus adapter (HBA) and firmware	Dual port, 4 GB HBAs (QLA2562)
Number of HBAs/host	2
Host servers	Intel Xeon CPU X7350 @ 2.93 GHz, 2.92 GHz, 4 Core, 128 MB memory
Total number of disks tested in the solution	320 for eight servers (per DB copy)
Maximum number of spindles that can be hosted in the storage	960

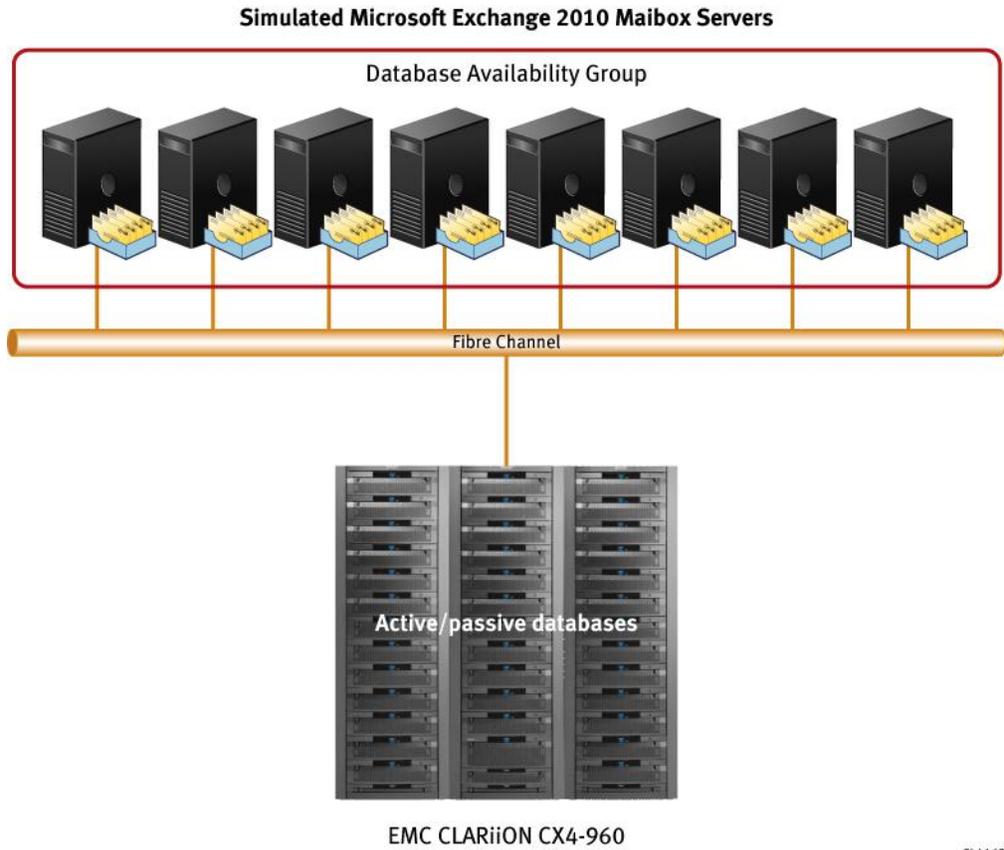
Storage software

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 5.3 (64 bit)
Host OS	Microsoft Windows Server 2008 Enterprise Service Pack 2
Exchange 2010 ESE.dll file	14.0.639.19

Physical environment

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



Exchange Server 2010 storage design

Overview The following sections provide an overview of the storage design guidelines followed when building this environment.

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

EMC can provide storage-optimized planning leveraging the EMC Storage Performance Sizer. Please contact your EMC Technical Consultant (TC) or e-mail Microsoft_Opportunities@emc.com for more information.

One of the methods that can be used to simplify the sizing and configuration of large amounts of EMC CLARiiON storage for use with Microsoft Exchange Server 2010 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 specific number of Exchange 2010 users. The amount of required resources is derived from a specific user profile type, mailbox size, and disk requirements.

Using the building-block approach takes out the guesswork and simplifies the implementation of the Exchange 2010 Mailbox server. 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.

Creating the building-block for the Exchange 2010 test environment

Now, understanding what a building-block is, administrators can create Exchange building-blocks based on the requirements provided in [Environment>Targeted customer profile](#).

The process of creating a building-block involves four simple steps:

1. Identify user requirements.
2. Identify and calculate storage requirements (based on both IOPS and capacity).
3. Identify your Exchange Mailbox server database design.
4. Finalize the Exchange Mailbox server storage configuration.

These four steps are described in greater detail in:

- [Exchange Server 2010 storage design> Step 1. Identify user requirements](#)
- [Exchange Server 2010 storage design>Step 2. Identify and calculate storage requirements](#)
- [Exchange Server 2010 storage design>Step 3. Identify the Exchange Mailbox server database design](#)
- [Exchange Server 2010 storage design>Step 4: Finalize Exchange Mailbox server storage configuration](#)

Step 1. Identify user requirements

The Exchange 2010 test environment needs to support a total of 80,000 users (with 10,000 users per server) at 0.15 IOPS per user, and a 1 GB MB mailbox quota.

Step 2. Identify and calculate storage requirements

Use this formula to calculate storage for the Exchange 2010 Mailbox server role:

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

Where	Is
IOPS	the number of input/output operations per second
%R	the percentage of I/Os that are reads
%W	the percentage of I/Os that are writes
WP	the RAID write penalty multiplier (RAID 1=2, RAID 5=4)
Physical Disk Speed	180 for 15k rpm drives or 140 for 10k rpm drives

Note

Microsoft also provides an Exchange 2010 Mailbox server role requirements calculator with some additional variables:

<http://msexchangeteam.com/files/12/attachments/entry453145.aspx>

Calculations:

The calculations for the test environment are summarized below. Calculations are based on the targeted user profile of .15 IOPS, and availability of 600 GB 10k rpm drives on the CLARiiON CX4-960.

Note

It is essential to calculate IOPS first, then capacity requirements.

IOPS Requirements:

- 10,000 users * .15 IOPS + 20% = 1,800 IOPS
- $(1,800 * .6) + 4(1,800 * .4) / 140 = 3960/140 = 28.2$ round-up to **30** spindles in RAID 5
- A total of **30** spindles are required to support IOPS requirements.

Capacity Requirements:

- Capacity for databases: 10,000 users * 1 GB mailbox + 35% = 13,500 GB
- Capacity requirements for logs: 10,000 users * 29 logs/user/day * 7 days deleted items retention = 2,030 GB
- Total space required based on capacity = **15,530 GB**
- To provide total required capacity for 15,530 GB it will require a total of **40** disks in eight RAID 5 (4+1) RAID groups on the CLARiiON CX4-960 ($15,530/2147 = 7.2$ round-up to 8)

Note

Five 600 GB disks in RAID 5 (4+1) will provide 2,147 GB of available formatted storage on the CX4-960.

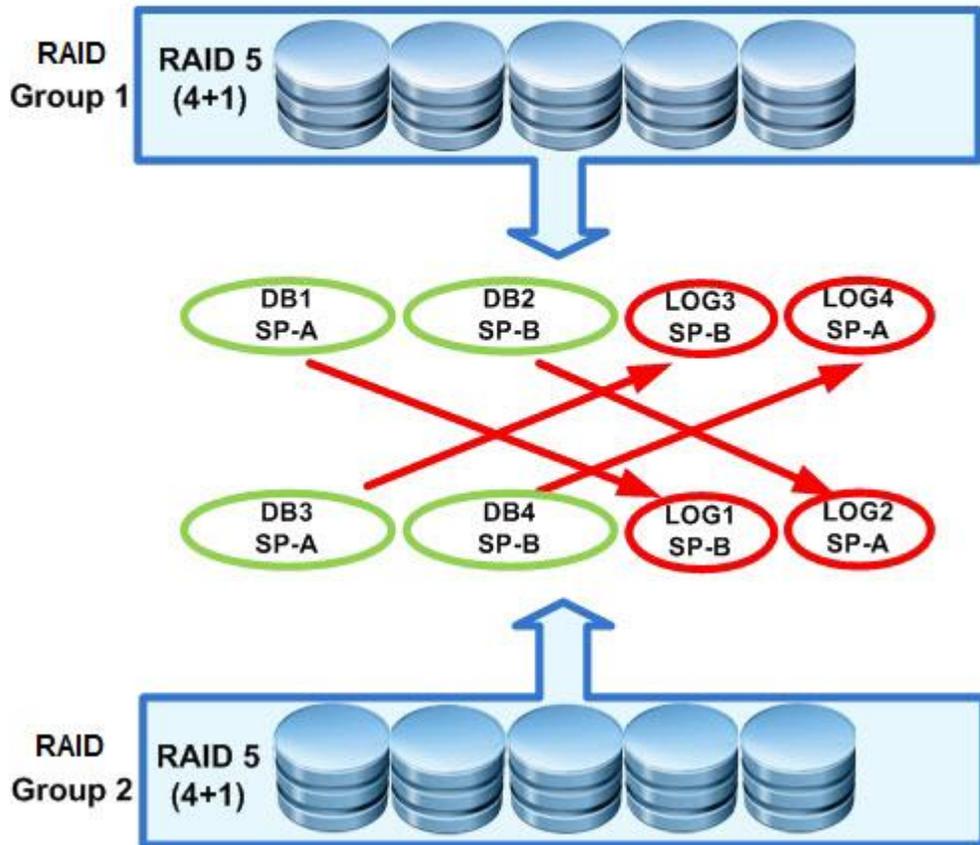
Number of Disks Required per DB copy:

Based on the above calculations, capacity requirements supersede IOPS requirements. Thus, a total of **40** disks per DB copy are needed to satisfy both requirements.

Step 3. Identify the Exchange Mailbox server database design

The next step is to identify how many databases to configure per Exchange server. This involves determining how large the databases need to be. The following image shows one of the suggested databases per RAID group design based on the RG capacity and Microsoft recommended guidelines for Exchange 2010.

CX4-960 Exchange 2010 storage Building-block design



Two Exchange databases and two log LUNs share five disks in a RAID 5 (4+1) configuration, but each Exchange database and its corresponding log LUNs are placed on separate RGs.

Based on the capacity of each RG, the database and log LUNs are configured to support 625 users, and will be able to accommodate two DBs and two log LUNs per RG. This totals 1,250 users per RG. In order to provide protection for logs, it is important to separate them from the same physical disks where the database resides. This would require a minimum of two RGs. Therefore, the minimum configuration with two RGs will support 2,500 users.

In summary, a building-block is created that provides all the necessary requirements for performance, capacity, and data protection to support 2,500 users.

Building block for the test configuration

The table below summarizes the final building-block created for this configuration based on the design approach illustrated in [Exchange Server 2010 storage design>Step 3: Identify your Exchange Mailbox server database design](#).

Item	Description
Number of users supported	2,500
User profile supported	0.15 IOPS—150 messages/user/day (30 sent/120 received)
Mailbox size	1 GB
Disk type and size	600 GB, 10k rpm FC
RAID type	RAID 5
Database LUN size	820 GB
Log LUN size	65 GB
Database/logs placement	Sharing spindles but in separate RAID groups (RGs)
Total disks required	Two RAID 5 RGs—10 disks (per database copy)

Step 4: Finalize the Exchange Mailbox server storage configuration

Scaling the configuration up to 10,000 users per the server requirement listed in [Environment>Targeted customer profile](#) will require four of these building-blocks to achieve the maximum configuration of 10,000 users. It will also require a total of 40 spindles with eight RAID 5 groups for each database copy.

Building-block scalability

Achieving the maximum configuration of 10,000 users (per the server requirement listed in [Environment>Targeted customer profile](#)) will require four building-blocks. This flexible design offers customers the capability to keep pace with an increasing user population. Users that share the same profile can easily be added to the environment, as shown in the following table.

Scaling to...	Requires...
2,500 users	10 spindles with two RAID 5 RGs
5,000 users	20 spindles with four RAID 5 RGs
7,500 user	30 spindles with six RAID 5 RGs
10,000 users	40 spindles with eight RAID 5 RGs

Note

If a second or third database copy is required, the total amount of spindles should be doubled or tripled based on the number of copies.

Exchange 2010 database maintenance

Database maintenance configuration options

Exchange 2010 database maintenance is now comprised of several tasks that maintain the integrity of mailbox databases. Database maintenance is divided into two parts:

- Store mailbox maintenance
- Extensible storage engine (ESE) database maintenance

In previous Exchange versions (Exchange 2007 SP1) ESE database maintenance was disk-intensive. In Exchange 2010, improvements have been made to increase performance in order to support both large mailboxes as well as to support different storage types.

Online defragmentation

The architecture for online defragmentation has changed in Exchange 2010. Online defragmentation was moved out of the mailbox database maintenance process, and now runs in the background 24x7. No settings need to be configured for this feature. Exchange monitors the database as it's being used, and small changes are made over time to keep it defragged for space and contiguity. Online defragmentation is also throttled so it doesn't have a negative impact on client performance. For more details on Exchange 2010 database maintenance, go to:

<http://technet.microsoft.com/en-us/library/bb125040.aspx>

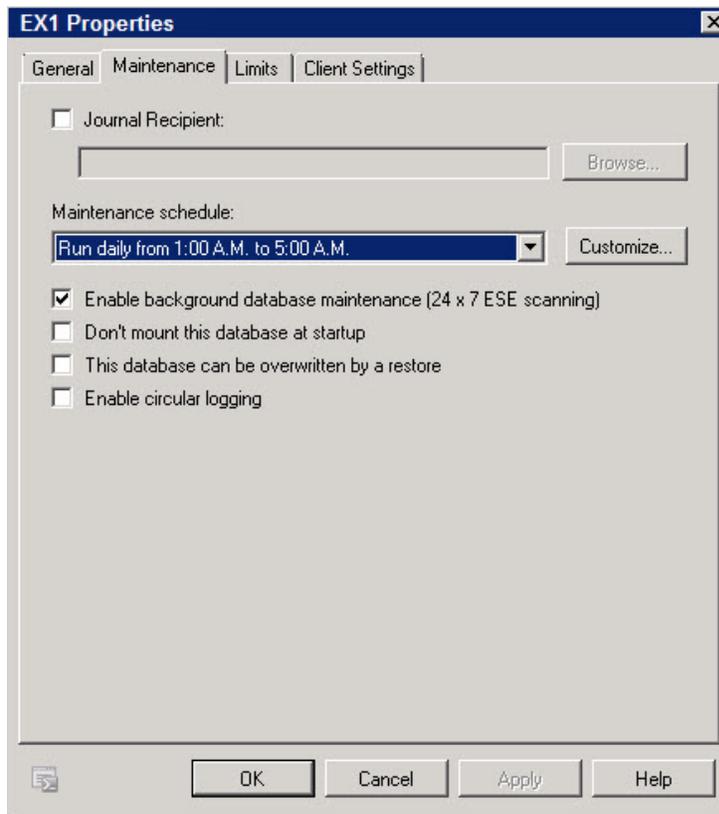
For more details on how to enable advanced ESE performance counters to monitor the online defragmentation process, go to:

<http://go.microsoft.com/fwlink/?LinkId=101194>

Online database scanning (or BDM)

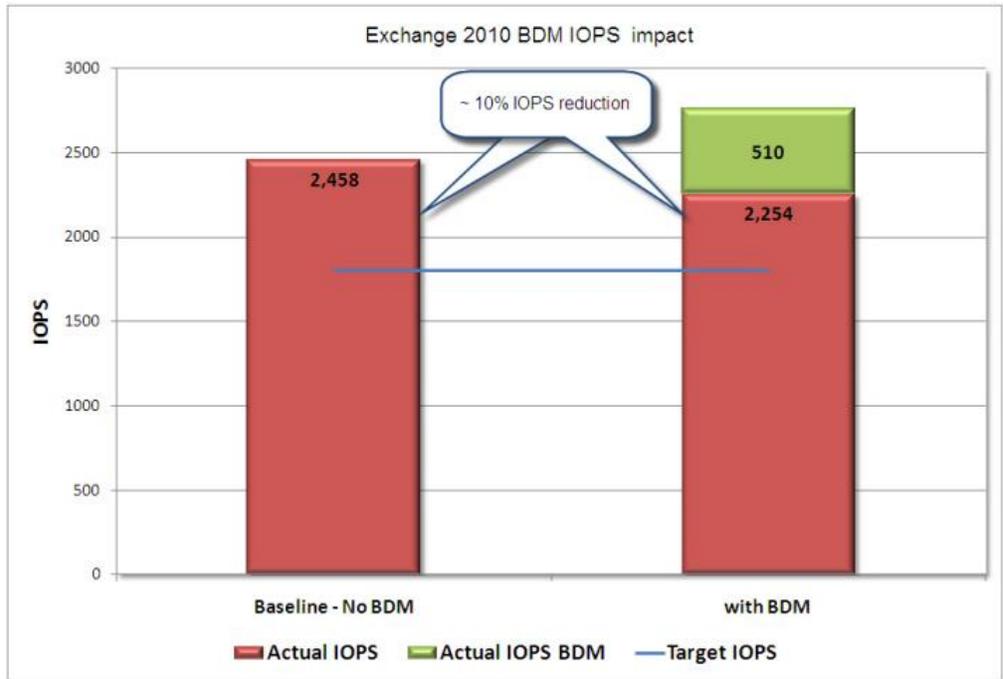
In Exchange 2010, there are now two options to run online database scanning (or BDM) on active database copies:

- Run BDM in the background 24x7 as the default behavior as shown in the image below (note, the **Enable background maintenance (24 x 7 ESE scanning)** checkbox is selected). This option is recommended for databases larger than 1 TB, where more time is needed to checksum the databases. Exchange scans the database no more than once per day, and again will generate a warning event if it can't finish scanning within a seven-day period.
- Run BDM by changing the Mailbox Database Maintenance schedule. This option is recommended for databases that are less than 1 TB in size. To enable this option, clear the **Enable background maintenance (24 x 7 ESE scanning)** checkbox.



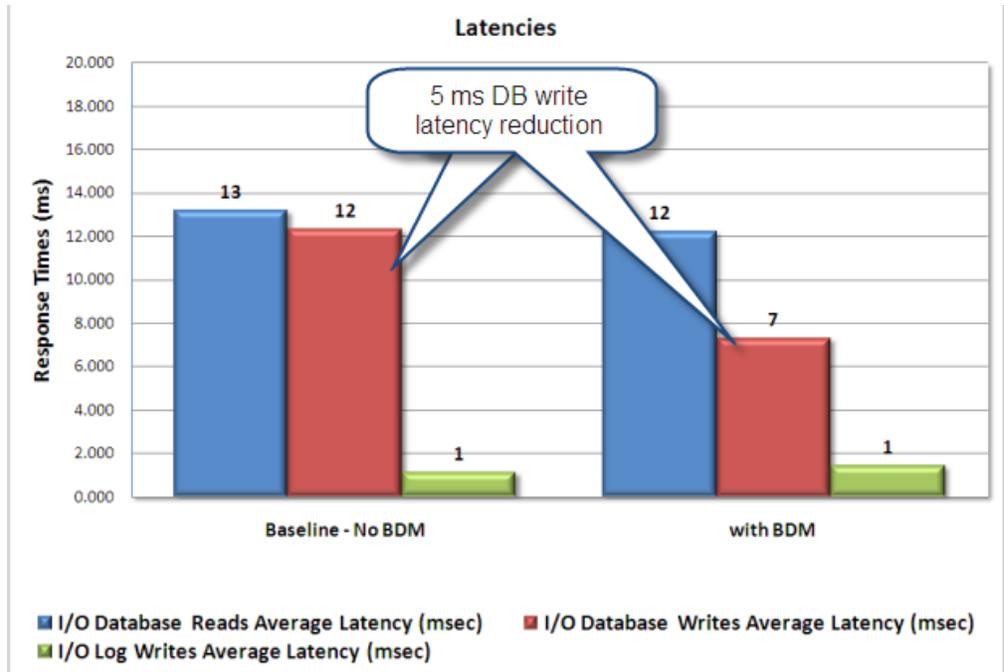
BDM effect on Exchange 2010 IOPS

Single Exchange 2010 server testing (10,000 users @ .18 IOPS) with and without BDM shows an approximate 10 percent IOPS reduction when BDM is enabled, as illustrated below.



BDM effect on Exchange 2010 latencies

Single Exchange 2010 server testing (10,000 users @ .15 IOPS) with and without BDM shows an approximate 5 ms DB write latencies decrease due to improved DB contiguity.



Best practices planning

Overview

In comparison to earlier versions of Microsoft Exchange, Exchange Server 2010 has made significant improvements in the areas of I/O and storage. For example, there have been many changes to the core schema, and the ESE to reduce the I/O usage profile.

Due to this I/O reduction, Exchange 2010 now supports SATA and SAS disks as well as FC drives and EFDs. These changes also allow for the use of RAID 5 as an optimal RAID configuration for Exchange databases and logs and even allow for RAID 0 (RAIDless) configurations under certain conditions. This allows for larger mailboxes at a reduced cost without performance degradation.

For Exchange 2010 Mailbox server design best practices, go to:

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

In addition to Microsoft recommendations, EMC recommends following the best practices described in this section to improve CLARiiON storage performance with Exchange 2010.

Mailbox servers

Follow these recommendations to ensure the best possible Mailbox server performance:

- Partition alignment is no longer required when running Microsoft Windows Server 2008 as partitions are automatically aligned to a 1 MB offset.

Note

Exchange Server 2010 requires Windows Server 2008.

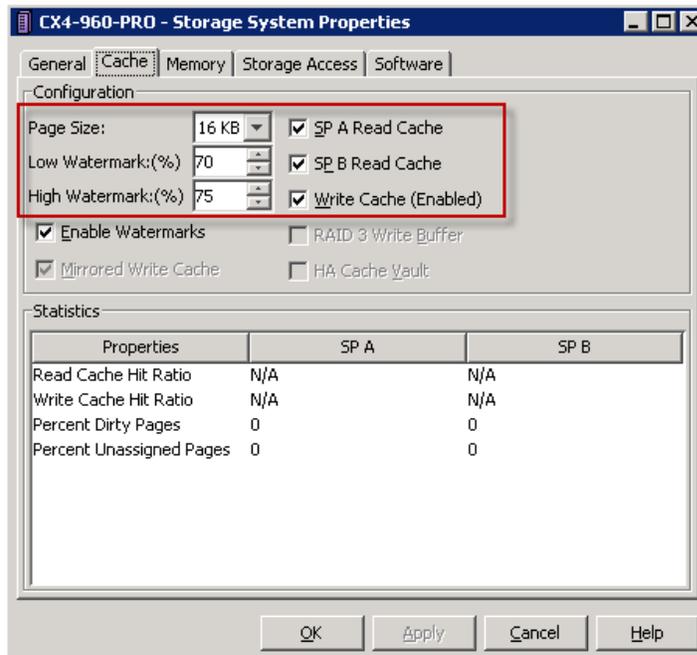
- When formatting a new NTFS volume to be used for Exchange databases and logs, it is recommended to set the allocation unit size (ALU) to 64 KB. This can be done from Disk Manager or through the CLI using `diskpart`.
- In SAN environments, use redundant HBAs connected to different fabrics.
- Install EMC PowerPath for optimal path management and maximum I/O performance. For more information on installing and configuring the PowerPath application, go to:

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

Tuning the CX4-960 storage parameters

Tuning the CX4-960 storage system parameters is important to ensure optimal performance. The table and image below detail the optimal configuration settings for a large enterprise deployment of Exchange 2010 users on an EMC CX4-960 storage array.

Parameter	Recommended setting
Page size	Set as follows: <ul style="list-style-type: none"> • 8 KB (if other applications share this array) • 16 KB (if the array is dedicated to an Exchange 2010 workload)
Low Watermark (%)	70
High Watermark (%)	75
Enable Watermark	Enabled
Write cache (Enabled)	Enabled



Requirements for placing logs and databases on the same physical spindles

In Exchange 2010, logs and databases can be placed on the same physical spindles under the following specific conditions:

- DAG is in place.
- DAG contains a minimum of three database copies.

When logs and databases cannot reside on the same physical spindles

When DAG is not being used, logs and databases should not share the same physical spindles. The exception to this is if there are multiple RGs in the storage configuration. Under these conditions, logs and databases may share the same spindles, but logs for the same database should be placed on different spindles.

For example, if there are two RGs being used for four databases (as shown in the [Exchange Server 2010 storage design>Step 3. Identify the Exchange Mailbox server database design](#) image) RAID Group 1 may contain:

- Database 1
- Database 2
- Logs 3
- Logs 4

RAID Group 2 may contain:

- Database 3
- Database 4
- Logs 1
- Logs 2

This structure provides additional protection to RAID when DAG is not used.

Storage sizing guidelines

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

- 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.
- Always size the Exchange environment for IOPS, then capacity.
- When calculating the IOPS requirements, always apply a 20 percent I/O overhead factor to your calculations to add some reserve (for antivirus, BDM, and so on).
- Balance storage disks across the array storage processors to take best advantage of CLARiiON performance and high availability.
- Microsoft recommends a maximum database size of 100 GB in environments where DAG is not being used. When DAG is being used with a minimum of two database copies, the maximum database size can be up to 2 TB. Consider backup (if applicable) and restore times when calculating the database size.

For more information on EMC solutions for Microsoft Exchange Server, go to:

<http://www.emc.com/solutions/application-environment/microsoft/solutions-for-microsoft-exchange-unified-communications.htm>

Testing methodology and results

Overview

Primary storage performance testing is designed to exercise the storage with maximum sustainable Exchange I/Os over a longer period of time. It is designed to show how long it takes for the storage to respond to an I/O under load.

Performance validation testing used the Jetstress tool to:

- Identify attainable IOPS for the CLARiiON CX4-960 (based on 80,000 Exchange users, using the building-block; see [Exchange Server 2010 storage design>Building block for the test configuration.](#))
 - Validate the total I/O throughput for eight Exchange servers.
 - Determine optimal CLARiiON SPs' utilization.
 - Measure volume shadow service (VSS) backup and log replay performance.
-

Test configuration

This section provides a high-level summary of the test data. Jetstress test results are provided based on the following configuration:

80,000 users, 8 servers (no BDM)

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 2010

Jetstress 2010 was used to verify the performance and stability of the CLARiiON CX4-960 disk subsystem. 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 specific 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.

For more information on Jetstress, go to:

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

Additional tools for assessing the environment

Although not applied per the scope of this validation testing, it is recommended to use the Microsoft Exchange Load Generator (LoadGen) tool for a full end-to-end assessment of the Exchange 2010 environment. Note that LoadGen can only be used in a non-production environment.

Additionally, Jetstress reports only application type (front-end) I/Os and does not analyze back-end storage. EMC CLARiiON Navisphere[®] Analyzer can be used to determine the total I/Os generated from each Exchange server to CX4-960 storage. EMC Navisphere Management enables you to manage, discover, monitor, and configure CLARiiON storage systems.

For more information on LoadGen and CLARiiON Navisphere Analyzer, go to the following websites:

<http://msexchangeteam.com/archive/2009/09/01/452272.aspx>

<http://www.emc.com/products/detail/software/navisphere-management-suite.htm>

Jetstress test types

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

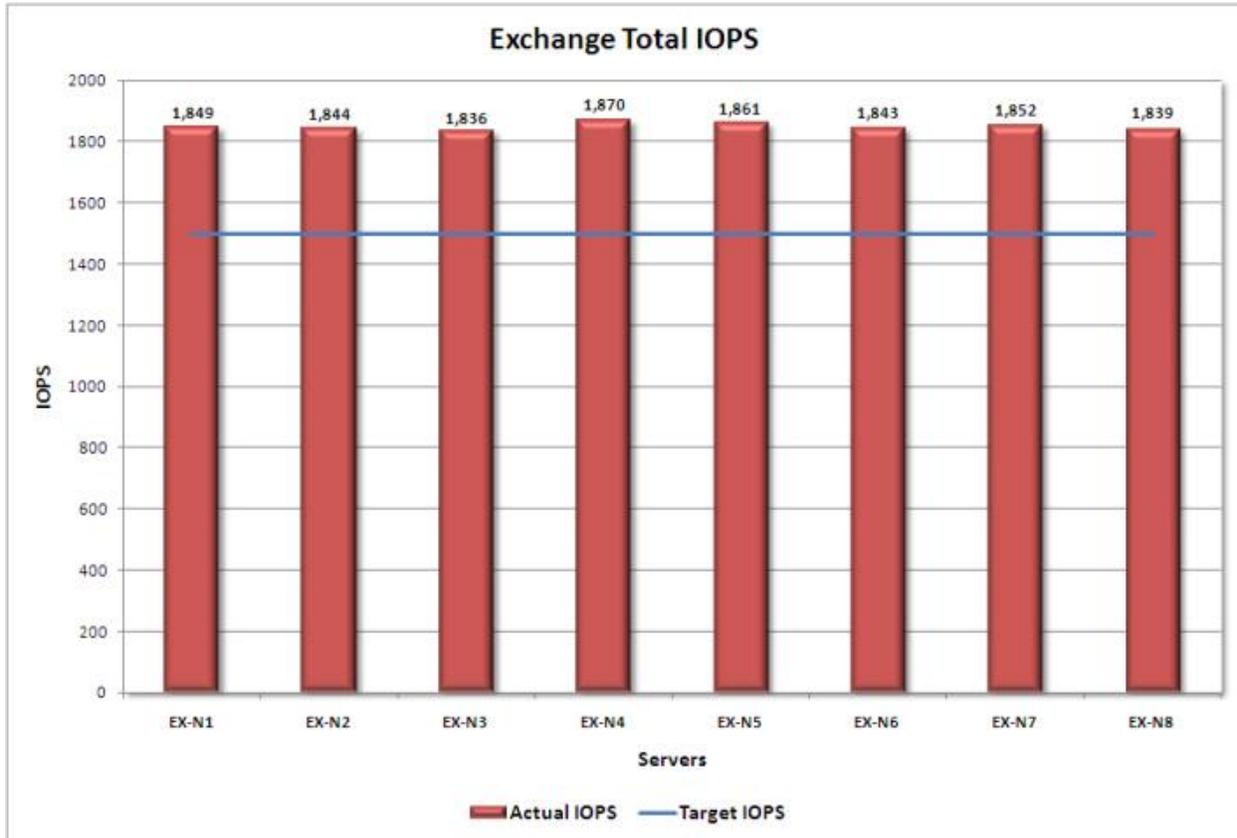
Test	Test name	Description
1	Two-hour Performance Test	Measures baseline storage I/O performance and how long it takes for the storage to respond to an I/O under load.
2	24-hour Stress Test	Validates how storage responds to a high I/O load for an extended period of time.
3	Database Backup VSS Test	Measures 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.

Individual server metrics

The following chart and table represent the total Jetstress IOPS results achieved on a per-server basis. Information includes:

- The sum of all of the achieved transactional IOPS
- The average of all the logical disks I/O latency in the 2-hour test duration

It is important to point out that the performance achieved is well within Microsoft's guidelines for Exchange 2010 subsystem performance, and achieved the target IOPS with extra headroom.



Item	Server 1	Server 2	Server 3	Server 4	Server 5	Server 6	Server 7	Server 8
Target Exchange transactional I/O per sec	1,800	1,800	1,800	1,800	1,800	1,800	1,800	1,800
Number of threads used	4	4	4	4	4	4	4	4
Achieved transactional I/O per sec	1,849	1,844	1,836	1,870	1,861	1,843	1,852	1,839
I/O database Reads/s	1,092	1,091	1,086	1,106	1,103	1,090	1,097	1,090
I/O database Writes/s	756	753	750	763	758	753	755	749
I/O database Reads average latency	16 ms	15 ms	16 ms					
I/O database Writes average latency	13 ms							
Total transaction log IOPS	596	595	593	608	608	596	598	599
Log replication I/O (I/O log Reads/s)	12	12	12	12	12	12	12	12
I/O log Writes/s	584	583	581	596	596	584	586	587
I/O log Writes average latency	3 ms							

Aggregate performance across all servers metrics

The sum of I/Os across servers and the average latency across all servers in the tested configuration is detailed in the table below.

The total IOPS achieved exceeded the target of 14,400 IOPS required to support 80,000 user configurations with a 0.18 IOPS tested profile.

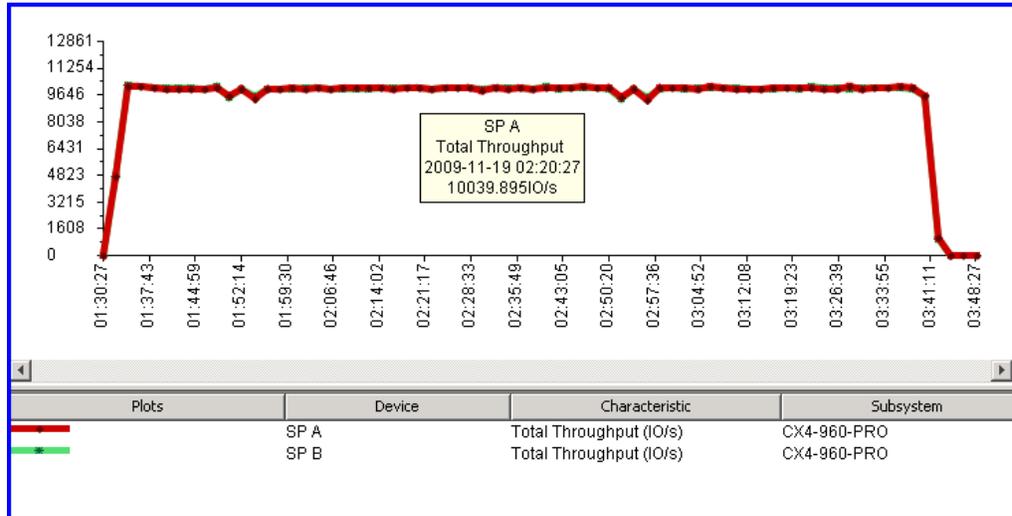
Database I/Os	For eight servers
Target Exchange transactional IOPS	14,400
Achieved database transactional IOPS	14,793
I/O database Reads/s	8,756
I/O database Writes/s	6,037
I/O database Reads average latency (ms)	16
I/O database Writes average latency (ms)	13
Log I/Os	
Total transaction log IOPS	4,793
Log replication I/O (I/O log Reads/s)	96
I/O log Writes/s	4,697
I/O log Writes average latency (ms)	3

Total throughput

The chart below shows the total I/O throughput for all eight Exchange servers, configured with 10,000 users at a 0.18 IOPS tested profile, while activity is generated against the CLARiiON CX4-960 array.

It shows the array handling about 20,000 IOPS (10,000 per SP). This translates to 2,500 total I/O activity generated from each Exchange server (20,000/8).

In comparison to the application front-end results from the table in [Testing methodology and results>Aggregate performance across all servers metrics](#), the CLARiiON back-end IOPS have increased by 20 percent because they include not only Jetstress application front-end IOPs, but also log replication IOPS.



Average SPs CPU utilization

While handling Exchange I/O activity, both CLARiiON SPs register 80 percent average CPU utilization. This means that the array is operating at optimal performance while handling additional Exchange IOPS introduced by database replication (two copies). If an additional load is required, a secondary CLARiiON array should be provisioned.

Database backup and recovery performance overview

Two Jetstress tests were performed 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.

MB read/s per database	13
MB read/s total per server	203

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.

Average time to play one log file	3 sec
--	-------

Conclusion

The goal of this white paper is to provide a simple framework that will guide customers in designing storage for Exchange 2010 solutions.

The building-block approach and methodologies used to deploy enterprise Exchange Server 2010 on the EMC CLARiiON CX4-960 storage array will help organizations to deploy Exchange 2010 more easily and efficiently.

Jetstress test results demonstrate that:

- A single CLARiiON CX4-960 array can handle the very heavy I/O activity of a large enterprise Exchange deployment.
- A single CLARiiON CX4-960 array is capable of handling two database copies for 80,000 users with a 1 GB mailbox, and a very heavy profile of .18 IOPS.

When BDM is enabled, the I/Os overhead factor should be included in the IOPS calculation requirements. This is to ensure that the same level of performance is achieved when BDM is enabled on the active database copies. For further details, see [Exchange 2010 BDM>BDM effect on Exchange 2010 latencies](#).

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