

EMC VNX5300 UNIFIED STORAGE 20,000 USERS WITH 1.5 GB MAILBOXES MICROSOFT EXCHANGE SERVER 2010 MAILBOX RESILIENCY STORAGE SOLUTION

Tested with: ESRP – Storage Version 3.0

Test date: July 2011

EMC SOLUTIONS GROUP

Abstract

This white paper describes the technical validation of a 20,000-user Exchange 2010 storage solution deployed on EMC VNX5300 Unified Storage according to criteria specified by the *Microsoft Exchange Solution Reviewed Program (ESRP) – Storage* program. The performance results and best practices presented in this paper provide validated guidelines for configuring the VNX5300 storage system for a large enterprise Exchange Server 2010 environment.

July 2011



Copyright © 2011 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 trademarks used herein are the property of their respective owners.

Part Number H8849

Table of contents

Overview	4
Disclaimer	4
Features	4
EMC VNX family of unified storage platforms.....	4
Software suites available	6
Software packs available.....	6
EMC VNX5300	6
Solution description.....	8
Storage design for Exchange Server 2010	9
Building block used in this solution	10
VNX storage pools.....	11
Additional factors that can affect mailbox server scalability.....	12
Targeted customer profile	12
Tested deployment	13
Best Practices	15
Mailbox servers.....	16
Networking	16
Core storage.....	17
Backup strategy	17
Information resources	18
Test results summary	18
Reliability (stress)	18
Storage performance.....	19
Individual server metrics	19
Performance across servers.....	19
Database backup and recovery performance.....	20
Database read-only performance.....	20
Transaction log recovery/replay performance	20
Detailed test results	21
How to view Jetstress reports	21
Conclusion	23
Contact EMC	23

Overview

This document provides information about EMC® VNX5300™ Unified Storage array performance with 20,000 Microsoft Exchange Server 2010 users, based on Microsoft Exchange Solution Reviewed Program (ESRP)—Storage program guidelines¹. For any questions or comments regarding the content of this document, see [Contact EMC](#) on page 23.

Disclaimer

This document has been produced independently of Microsoft Corporation. Microsoft Corporation expressly disclaims responsibility for, and makes no warranty—express or implied—with respect to the accuracy of the contents of this document.

The information contained in this document represents the current view of EMC on the issues discussed as of the date of publication. Due to changing market conditions, it should not be interpreted as a commitment on the part of EMC. In addition, EMC cannot guarantee the accuracy of any information presented after the date of publication.

Features

EMC VNX family of unified storage platforms

The EMC VNX™ family delivers industry-leading innovation and enterprise capabilities for file, block, and object storage in a scalable, easy-to-use solution. This next-generation storage platform combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today's enterprises.

All of this is available in a choice of systems ranging from affordable entry-level solutions to high-performance, petabyte-capacity configurations servicing the most demanding application requirements. The VNX family includes the VNXe™ series, purpose-built for the IT manager in entry-level environments, and the VNX series, designed to meet the high-performance, high-scalability requirements of midsize and large enterprises.

¹ The ESRP—Storage program was developed by Microsoft Corporation to provide a common storage testing framework for vendors to provide information on their storage solutions for Microsoft Exchange Server software. For more details on the Microsoft ESRP—Storage program, refer to <http://technet.microsoft.com/en-us/exchange/ff182054.aspx>.



Figure 1. EMC VNX family of unified storage platforms

The VNX series delivers uncompromising scalability and flexibility for the mid-tier and enterprise space while providing market-leading simplicity and efficiency to minimize total cost of ownership. Customers can benefit from the new VNX features such as²:

- Next-generation unified storage, optimized for virtualized applications
- Extended cache using Flash drives with FAST Cache and Fully Automated Storage Tiering for Virtual Pools (FAST VP) that can be optimized for the highest system performance and lowest storage cost simultaneously on both block and file
- Multiprotocol support for file, block and object with object access through Atmos™ Virtual Edition (Atmos VE)
- Simplified management with EMC Unisphere™ for a single management framework for all NAS, SAN, and replication needs
- Up to three times improvement in performance with the latest Intel multicore CPUs, optimized for Flash
- 6 Gb/s SAS back end with the latest drive technologies supported:
 - 3.5” 100 GB and 200 GB Flash, 3.5” 300 GB, and 600 GB 15k or 10k rpm SAS, and 3.5” 2 TB 7.2k rpm NL-SAS
 - 2.5” 300 GB and 600 GB 10k rpm SAS
- Expanded EMC UltraFlex™ I/O connectivity—Fibre Channel (FC), Internet Small Computer System Interface (iSCSI), Common Internet File System (CIFS), Network File System (NFS) including parallel NFS (pNFS), Multi-Path File System (MPFS), and Fibre Channel over Ethernet (FCoE) connectivity for converged networking over Ethernet

The VNX series includes five new software suites and three new software packs, making it easier and simpler to attain the maximum overall benefits.

² Features listed are based on the VNX operating environment version available at the time of this solution validation. EMC constantly improves and updates its storage fleet with new features and functionalities. For latest features and updates visit www.emc.com.

Software suites available

- VNX FAST Suite—Automatically optimizes for the highest system performance and the lowest storage cost simultaneously (FAST VP is not part of the FAST Suite for the VNX5100™).
- VNX Local Protection Suite—Practices safe data protection and repurposing.
- VNX Remote Protection Suite—Protects data against localized failures, outages and disasters.
- VNX Application Protection Suite—Automates application copies and proves compliance.
- VNX Security and Compliance Suite—Keeps data safe from changes, deletions, and malicious activity.

Software packs available

- VNX Total Efficiency Pack—Includes all five software suites (not available for the VNX5100).
- VNX Total Protection Pack—Includes local, remote, and application protection suites.
- VNX Total Value Pack—Includes all three protection software suites and the Security and Compliance Suite (the VNX5100 exclusively supports this package).

For additional details about the EMC VNX family of unified storage systems, refer to the white paper *Introduction to EMC VNX series* at

<http://www.emc.com/collateral/hardware/white-papers/h8217-introduction-vnx-wp.pdf>.

EMC VNX5300

The EMC VNX5300 model is a member of the VNX series next-generation storage platform, powered by the Intel four-core Xeon 5600 series processors with a 6 Gb SAS drive back end, providing the industry's highest bandwidth. The VNX5300 model provides high-performing, unified storage with unsurpassed simplicity and efficiency. Organizations will achieve new levels of performance, protection, compliance, and ease-of-management.

The VNX5300 storage array delivers a single-box block and file solution, which offers a centralized point of management for distributed environments. This makes it possible to dynamically grow, share, and cost-effectively manage multiprotocol file systems and provide multiprotocol block access.

Table 1 lists VNX5300 features for block. For additional VNX specifications for both block and file, visit www.emc.com at:

<http://www.emc.com/collateral/software/specification-sheet/h8514-vnx-series-ss.pdf>

Table 1. VNX5300 features summary (shown for block only)

System feature	Value
Minimum/Maximum drives per system	4/125
Drive Enclosure Options (DAE)	25 x 2.5" SAS/Flash drives–2U 15 x 3.5" SAS/Flash drives–3U
Array enclosure	3U disk processor enclosure (Holds 15x3.5" or 25x2.5" SAS/Flash drives)
Drive types	Flash, SAS, NL-SAS
Drive enclosure options	25 x 2.5" SAS/Flash drives – 2U 15 x 3.5" SAS/Flash drives – 3U
I/O architecture	PCI-e Gen 2
RAID options	0 / 1 / 10 / 3 / 5 / 6
CPU / Memory per array	Quad-core Intel Xeon 5600 @ 1.6 GHz / 16 GB (8 GB per SP)
File protocols	NFS, CIFS, MPFS, pNFS
Blocks (Number of SPs)	2
Block protocols	FC, iSCSI, FCoE
Max block flex I/O modules per array	4
Max raw capacity	240 TB
Max SAN hosts	2,048
Max number of pools	20
Max number of LUNs	2,048
Max total ports per array	24
2/4/8 Gb/s FC Max ports per array	16
1 GBaseT iSCSI Max total ports per array	8
10 GbE iSCSI Min/Max total ports per array	4
FCoE Max total ports per array	4
6 Gb/s SAS buses for DAE connections	2
Management and base software	Integrated software includes: <ul style="list-style-type: none"> • Unisphere • File deduplication/compression • Block compression • Virtual Provisioning™ • SAN Copy

Solution description

This solution is intended for medium to large enterprise-size businesses that are planning to deploy Microsoft Exchange Server 2010 on EMC storage and would like to leverage the Exchange Server 2010 mailbox resiliency Database Availability Group (DAG) feature. The solution design represents an Exchange Server 2010 environment supporting 20,000 users in a mailbox resiliency configuration across two intelligent VNX5300 storage arrays. The solution is designed to provide outstanding performance and flexibility for today's and tomorrow's Exchange users.

In this solution, 20,000 users are deployed across a DAG with four Exchange mailbox servers, two servers in the primary site and two in the secondary site. Each database has two DAG copies, one RAID-protected by a VNX5300 array and the other replicated to a different server on a secondary array. Each Exchange mailbox server hosts 5,000 active users and 5,000 passive users during normal operation, and supports 10,000 active users should a database switchover from another mailbox server due to a server failure or for maintenance purposes.

The user mailboxes have a 1.5 GB capacity and 0.10 IOPS profile with a 20 percent throughput reserve, also known as "overhead" (so 0.12 IOPS per user is tested in Jetstress).

The solution is designed to eliminate a single point of failure and handle the loss of an array, mailbox server, database volume, HBA, or switch.

Note: This solution architecture can be compared to one in which Exchange Server 2010 is deployed in a stand-alone configuration (no DAGs) and a single array with two mailbox servers provides service to all 20,000 users.

For more information, access the Windows Hardware Compatibility List link for the EMC VNX5300 unified storage array at

<http://www.windowsservercatalog.com/item.aspx?itemId=6d5fa064-ca7f-bfb5-9e76-f4c39dc1c667&bCatID=1282>

Figure 2 illustrates the architecture of the EMC VNX5300 20,000 Mailbox Exchange Server 2010 mailbox resiliency storage solution.

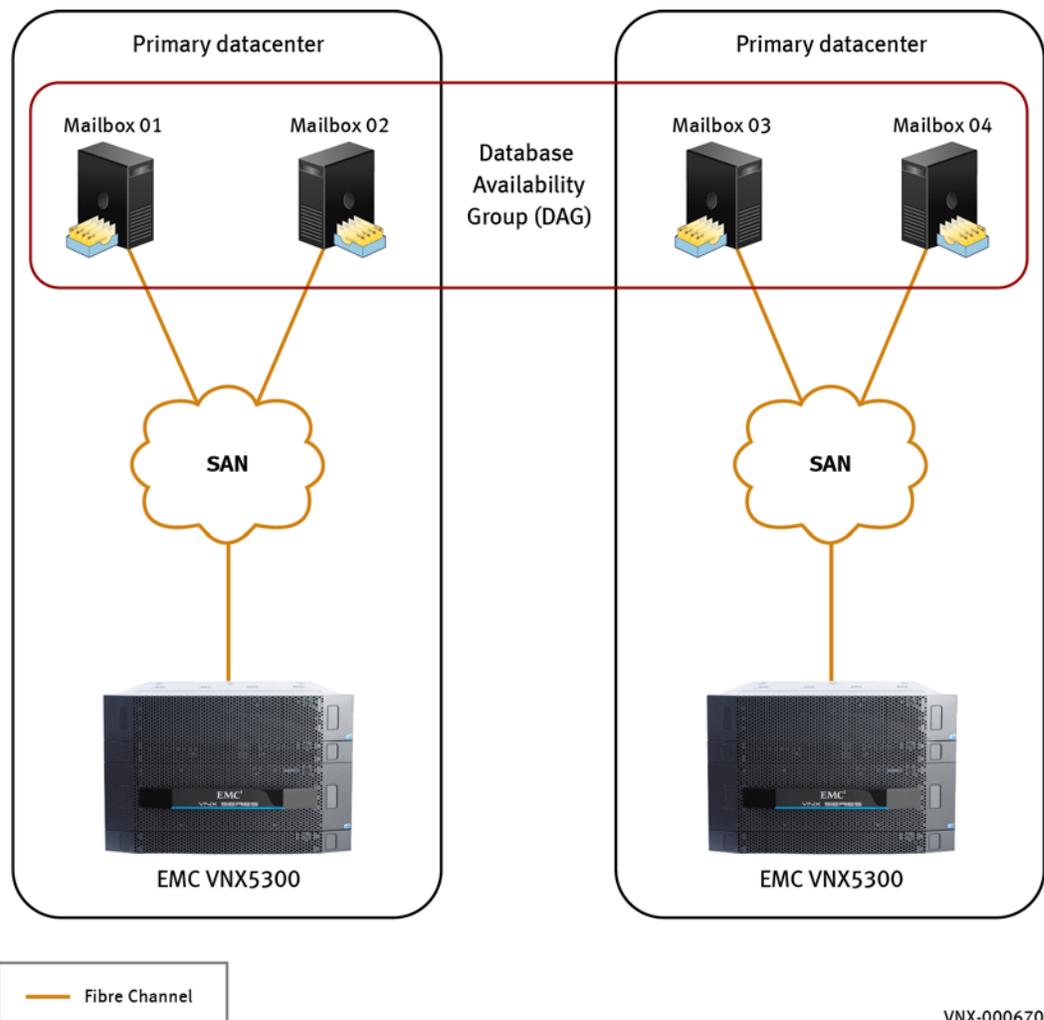


Figure 2. Solution architecture

Storage design for Exchange Server 2010

Storage design is an important element for ensuring the successful deployment of Microsoft Exchange Server 2010. Sizing and configuring storage for use with Exchange Server 2010 can be a complicated process, driven by many variables and requirements, which vary from organization to organization. Properly configured Exchange storage, combined with optimally sized server and network infrastructures, can guarantee a smooth Exchange operational environment.

One of the methods that can be used to simplify the sizing and configuration of large amounts of storage on EMC VNX series storage arrays for use with Exchange Server 2010 is to define a unit of measure—a mailbox server building block.

A mailbox server building block represents the amount of disk and server (CPU, memory, and network) resources required to support a specific number of Exchange Server 2010 users. The amount of required resources is derived from a specific user profile type, mailbox size, and high availability requirements. Using the building block approach simplifies the design and implementation of Exchange Server 2010.

Once the initial building block is designed, it can be easily reproduced to support the required number of users in your enterprise. By using this approach, Exchange administrators can create their own building-blocks based on their company's Exchange environment requirements. This approach is very helpful when a customer expects future growth, as it makes Exchange environment additions much easier and straightforward.

EMC's best practices involving the building-block approach for an Exchange Server design has been very successful for many customer implementations.

Building block used in this solution

In this solution, the mailbox server building block is made up of 34 2 TB NL-SAS drives to support 10,000 active users in a single Exchange mailbox server, with a 1.5 GB mailbox size and 0.10 IOPS per user (0.12 IOPS is tested to include a 20 percent IOPS reserve). We created a single RAID 1/0 storage pool that consists of 32 NL-SAS drives to house Exchange database files, and built a RAID 1/0 (1+1) RAID group of two NL-SAS drives for the Exchange mailbox server logs.

This solution uses two of these building blocks (68 NL-SAS drives in total) on the primary storage array to scale the configuration up to 20,000 users. The two building blocks are duplicated to the secondary storage array.

Table 2 summarizes the attributes of the building block.

Table 2. Mailbox server building block attributes

Item	Description
Number of users supported by a single mailbox server	10,000
User profile	0.10 IOPS with a 20% reserve (0.12 IOPS tested)
Mailbox size	1.5 GB
Drive type and capacity	2 TB NL SAS drives, 7,200 rpm
RAID type	RAID 1/0
Database LUN size	1.3 TB
Log LUN size	50 GB
Number of disks to support a single mailbox server	34

VNX storage pools

A storage pool is a single repository of physical disks on which logical units (LUNs) are created. Pools can contain a few disks or hundreds of disks, whereas RAID groups are limited to 16 disks. Because of the large number of disks supported in a pool, pool-based provisioning provides similar benefits of metaLUN striping across many drives; however, unlike metaLUNs, it requires minimal planning and management effort. Pools support a single RAID protection level that can be RAID 5, RAID 6, or RAID 1/0. The use of storage pools simplifies storage provisioning.

Note: Single RAID level protection within the pool is based on the current release of the VNX operating environment. Future releases will provide the capabilities to deploy multiple RAID levels within same storage pool. Periodically review VNX features on www.emc.com.

Figure 3 illustrates the building block's database and log LUN configuration in this solution, designed to provide the best performance and sufficient capacity for 10,000 users with a 1.5 GB mailbox size and 0.10 IOPS per user (0.12 IOPS tested to include overhead). Twenty database LUNs are created from each RAID 1/0 storage pool, and twenty log LUNs are created from each RAID 1/0 (1+1) RAID group.

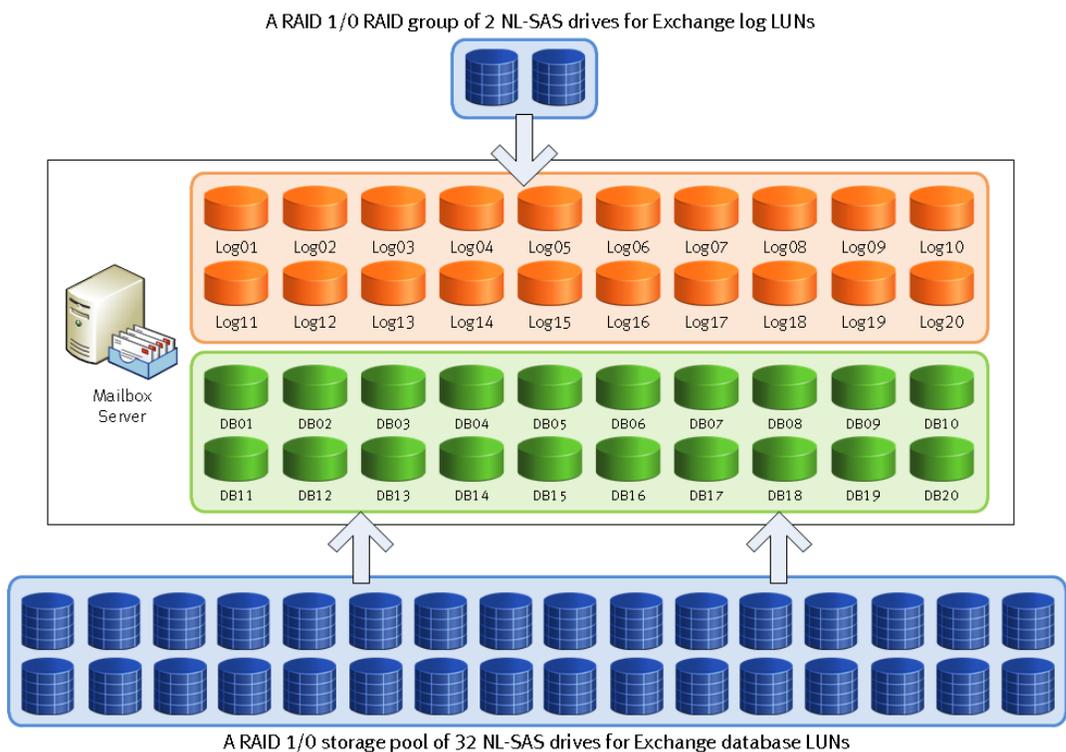


Figure 3. Building block's database and log LUN configuration

Additional factors that can affect mailbox server scalability

The ESRP—Storage program focuses on storage solution testing to address performance and reliability issues with storage design. However, storage is not the only factor to consider when designing a scalable Exchange solution. Other factors that can affect server scalability include:

- Server processor utilization
- Server physical and virtual memory limitations
- Resource requirements for other applications
- Directory and network service latencies
- Network infrastructure limitations
- Replication and recovery requirements
- Client usage profiles

Any combination of these factors can affect the total number of mailboxes supported by a single mailbox server. Since all of these factors fall outside the scope of ESRP—Storage program solution validation, a single mailbox server might not necessarily support 10,000 users (as validated for the solution described in this document) in every customer deployment.

For more information on identifying and addressing performance bottlenecks in an Exchange system, refer to Microsoft's Performance and Scalability guide, available at <http://technet.microsoft.com/en-us/library/dd351197.aspx>.

Microsoft supports Exchange Server 2010 in production on hardware virtualization software such as Hyper-V and VMware vSphere and any third-party hypervisor that has been validated under the Windows Server Virtualization Validation Program. If Exchange mailbox servers are deployed as virtual machines (VMs), the number of users that can be supported by each mailbox server VM depends on the hypervisor type, physical server hardware capabilities such as CPU and memory characteristics, and high availability requirements. For more details about Exchange Server 2010 system requirements for hardware virtualization, visit <http://technet.microsoft.com/en-us/library/aa996719.aspx>.

Targeted customer profile

This solution is designed for any medium-to-large enterprise that is planning to consolidate its Exchange Server 2010 storage environment and requires a highly reliable, scalable storage solution. The design of this solution has the following characteristics:

- 20,000 active mailboxes
- 100 messages sent/received per user per day, 0.10 IOPS throughput with a 20 percent overhead (0.12 IOPS tested)
- 1.5 GB mailbox size
- Two Exchange servers with a 5,000 active/5,000 passive configuration (tested by simulating 10,000 active users per server)
- Native DAG replication mechanism for mailbox resiliency and high availability with two database copies maintained

- 24x7 background database maintenance (BDM) enabled
- Two EMC VNX5300 storage arrays (the solution is tested with two mailbox servers connected to a single array)

Note: If three database copies are required, EMC recommends placing two RAID-protected DAG copies on a local array and placing a third RAID-protected copy on another array.

Tested deployment The following tables summarize the testing environment.

Table 3. Simulated Exchange configuration

Attribute	Value
Number of Exchange mailboxes	20,000
Number of Database Availability Groups (DAGs)	1
Number of servers/DAG	4 (2 tested)
Number of active mailboxes/server	5,000 (10,000 simulated in Jetstress testing)
Number of databases/mailbox server	20
Number of copies/database	2
Number of mailboxes/database	500
User profile	0.10 (0.12 tested in Jetstress)
Database LUN size	1.3 TB
Log LUN size	50 GB
Total database size for performance testing	40 TB (20 TB for each server)
Percentage of storage capacity used by Exchange database ³	77% (40/52)

Table 4. Storage hardware

Component	Description
Storage Connectivity (FC or iSCSI)	FC
Storage model and OS/firmware revision	VNX5300, VNX Block Operation Environment version: 05.31.000.5.011 http://www.windowsservercatalog.com/item.aspx?idItem=6d5fa064-ca7f-bfb5-9e76-f4c39dc1c667&bCatID=1282
Storage cache	8 GB mirrored
Number of storage controllers	Two storage processors (SPs)

³ Storage performance characteristics change based on the percentage utilization of the individual disks. Tests that use a small percentage of the storage (~25 percent) may exhibit reduced throughput if the storage capacity utilization is significantly increased beyond what is reported in this paper.

Component	Description
Number of storage ports	4 (Two for each storage processor)
Maximum bandwidth of storage connectivity to host	16 Gbps (2*8 Gbps FC Ports)
Switch type/model/firmware revision	Cisco MDS 9509 FC switch, 4 Gbps, Firmware 3.2.2c
HBA model and firmware	Emulex LPe12002-E 8Gb 2-port PCIe Fibre Channel Adapter, Firmware 1.00A12
Number of HBAs/host	2
Host server type	Dell PowerEdge R810 with Intel(R) Xeon(R) L7555, 1.87 GHz CPU, Eight cores, 128 GB RAM
Number of disks used in solution	68
Maximum number of disks supported by VNX5300 storage array	125

Table 5. Storage software

Component	Description
HBA driver	7.2.41.2
HBA QueueTarget Setting	0
HBA QueueDepth Setting	32
Multi-pathing	EMC PowerPath 5.5 SP1
Host OS	Microsoft Windows Server 2008 R2 Enterprise SP1
ESE.dll file version	14.01.0279.000
Replication solution name/version	N/A

Table 6. Storage disk configuration (mailbox store disks)

Attribute	Description
Disk type, speed and firmware revision	2 TB NL SAS 7.2k RPM, Firmware BS17
Raw capacity for each disk (GB)	1,834 GB
Number of physical disks in test	64
Total raw storage capacity (GB)	117, 376 GB
RAID level	RAID 1/0
Total formatted capacity	52 TB (26 TB for each server)
Storage capacity utilization (percentage)	45% (52*1024/117, 376)
Database capacity utilization (percentage)	35% (40*1024/117,376)

Table 7. Storage disk configuration (transactional log disks)

Attribute	Description
Disk type, speed, and firmware revision	2 TB NL SAS 7.2k RPM, Firmware BS17
Raw capacity for each disk (GB)	1,834 GB
Number of physical disks in test	4
Total raw storage capacity (GB)	7,336 GB
RAID level	RAID 1/0
Total formatted capacity	2 TB (1 TB for each server)

Table 8. Replication Configuration

Attribute	Description
Replication mechanism	Exchange 2010 DAG Mailbox Resiliency
Number of links	2
Simulated link distance	LAN
Link type	IP
Link bandwidth	Gigabit Ethernet (1 Gbps)

Best Practices

Microsoft Exchange Server 2010 has changed significantly since earlier versions of Exchange, particularly with regard to I/O and storage. Exchange 2010 is designed to reduce the amount of I/O traffic to a storage subsystem than was necessary in previous Exchange versions. Microsoft accomplished this by implementing a larger page size (32 K) and more effectively using mailbox server cache through the use of larger log checkpoint depth (100 MB in Mailbox Resiliency configurations). When the database has a higher checkpoint depth target, the system is able to retain database file changes in memory for a longer period, thus improving the ability to consolidate I/O activity (coalescence). When more I/Os are coalesced, repeated write I/Os are reduced by being delayed long enough so that multiple database changes can be made in memory prior to writing the changes to the database file.

For more information about the Exchange 2010 new storage features, visit <http://technet.microsoft.com/en-us/library/bb125040.aspx>

Due to this I/O reduction, customers can now deploy Exchange Server 2010 on low cost SATA and NL SAS disks in addition to FC, SAS, and SSD drives. Depending on the Exchange Server 2010 deployment model the customer chooses, DAG or stand-alone, storage configuration options can vary from highly reliable SAN infrastructure to Direct Attached storage (DAS), to even non-RAID JBOD configurations in specific circumstances. All RAID types (RAID 0, RAID 5, RAID 1/0 and RAID 6) are supported to meet the requirements of the high availability model and storage configuration you choose to deploy. EMC provides storage options for all of these configurations.

- For Exchange Server 2010 storage design best practices, visit <http://technet.microsoft.com/en-us/library/ee832792.aspx>
- For Exchange Server 2010 mailbox server design best practices, visit <http://technet.microsoft.com/en-us/library/dd346703.aspx>

In addition to following Microsoft's recommendations, EMC recommends that you follow the best practices described in this section to improve EMC VNX series storage performance with Exchange Server 2010.

Mailbox servers

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

- Format NTFS volumes (Windows 2008 or Windows 2008 R2) to be used for Exchange databases and logs with an Allocation Unit size of 64 KB.
Note: Partition alignment is no longer required when running Microsoft Windows Server 2008 or Windows 2008 R2 since partitions are automatically set to a 1 MB alignment offset.
- In SAN environments, use redundant host bus adapters (HBAs) connected to different fabrics.
- Verify that the HBA installed in the server can support the IOPS requirements, even in a failover situation. To avoid throttling, ensure that the queue depth is set according to EMC recommendations for the specific HBA vendor and type.
- Apply all required OS hotfixes recommended by Microsoft, EMC, and the HBA vendor.
- Install EMC PowerPath® for optimal path management and maximum I/O performance. For more information on installing and configuring EMC PowerPath, visit <http://www.emc.com/products/detail/software/powerpath-multipathing.htm>.

Networking

For high availability deployments that use DAGs, multiple physical NICs connected to different networks are recommended to isolate user MAPI traffic from database replication traffic.

For iSCSI deployments, multiple network switches are preferred for fault tolerance and performance. Where this is not possible, VLANs must be used to isolate iSCSI traffic from all other network traffic.

Core storage

EMC VNX5300 is already optimized for Exchange Server 2010 workloads and requires only a few minor adjustments. Follow these guidelines to ensure optimal storage performance:

- Always calculate the I/O spindle requirements first, and then calculate the capacity requirements.
- After establishing actual required IOPS, apply a 20 percent IO “overhead” factor to build in adequate reserve throughput capability.
- Plan for performance even in a failover situation.
- Consider the additional bandwidth requirements imposed by BDM.
- Set the VNX storage array page size to 16 K.
- Ensure that write cache is enabled for all database and log LUNs.
- Isolate the Exchange server database workload to a different set of spindles from other I/O-intensive applications or workloads (assign the Exchange workload to its own set of disks). This ensures the highest level of performance for Exchange and simplifies troubleshooting in the event of a disk-related Exchange performance issue.
- Both storage pools and RAID groups work well with Exchange Server 2010.
- When using storage pools for Exchange data:
 - Use homogeneous storage pools (storage pools made up of disks of the same type) with the appropriate RAID multiplier: 8 (4+4) drives for RAID 1/0 pools, 5 (4+1) drives for RAID 5 pools, and 8 (6+2) drives for RAID 6 pools.
 - Do not use the FAST VP feature.
 - Isolate each DAG copy in its own storage pool.
 - For optimal performance, isolate databases from logs on separate sets of disks.
 - Balance LUNs across the array storage processors to take the most advantage of VNX5300 performance and high availability, and distribute the I/O load evenly across VNX5300 front-end ports and back-end buses for failover and load balancing.
 - Microsoft recommends a maximum database size of 200 GB in environments that do not use DAG. When DAG is 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.

Backup strategy

When using intelligent storage, use VSS snapshots or clones to back up and protect your Exchange data.

Information resources

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

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

Test results summary

This section provides a high-level summary of the ESRP test results for this solution.

Note: Detailed test result reports are attached to this PDF in HTML format.

The solution was validated using Microsoft Jetstress 2010 to ensure that the storage design satisfies the disk I/O and capacity requirements for the target profile. Jetstress simulates Exchange Server 2010 I/O types at the database level by interacting directly with the database technology of the ESE (also known as Jet) on which Exchange is built.

Jetstress can be configured to test the maximum I/O throughput available to the disk subsystem within the performance constraints of Exchange. Alternatively, Jetstress can be configured to accept a specific user profile (user count, IOPS per user) and validate that the disk subsystem is capable of maintaining an acceptable performance level with such a profile. Both configuration options produce similar results.

Note: Jetstress is designed to test the performance of an Exchange storage subsystem before you move it into a production environment. Jetstress is not designed to test server CPU, memory configuration, or the impact of MAPI user activity. To test Exchange server functionality and end-to-end deployment, it is a best practice to use Microsoft Load Generator (Loadgen). For more information about the Loadgen tool, visit <http://www.microsoft.com/downloads/en/details.aspx?FamilyID=cf464be7-7e52-48cd-b852-ccfc915b29ef>

The documentation for Jetstress describes how to configure and execute IO validation or evaluation on your server hardware. The Microsoft Jetstress application is available at <http://go.microsoft.com/fwlink/?LinkId=178616>

Reliability (stress)

The reliability (stress) test runs for 24 hours. The goal is to validate that the storage can handle high I/O load for a long period of time. Both log and database files are analyzed for integrity after the stress tests to ensure that there is no database or log corruption. The 24-hour stress test results reveal:

- No errors in the saved event log file.
- No errors for the database and log checksum processes.

Storage performance

The primary storage performance testing is designed to exercise the storage with maximum sustainable Exchange I/O activity for two hours. The test shows how long it takes for the storage to respond to an I/O under load. The data below summarizes all of the logical disk I/Os and average of all the logical disks I/O latency in the two-hour test duration. Each server is listed separately and the aggregate numbers across all servers is listed as well.

Individual server metrics

Table 9 presents the sum of transactional I/Os across all databases (database read operations per second and database write operations per second) and the average latencies across all databases and logs for each server. The configuration is designed to achieve a target of 1,200 IOPS for 10,000 users on each server. This includes 20 percent overhead, above the 0.10 IOPS user profile, for reserve throughput capability.

Table 9. Individual server metrics

Database I/O	MBX01	MBX02
Achieved transactional I/Os	1340.745	1344.478
I/O database reads/sec	830.439	832.943
I/O database writes/sec	510.306	511.531
I/O database average read latency (ms)	17.498	17.558
I/O database average write latency (ms)	3.051	2.994
Transaction log I/O	MBX01	MBX02
I/O log writes/sec	468.626	468.599
I/O log average write latency (ms)	1.175	1.182

Performance across servers

Table 10 presents the sum of transactional I/Os and the average latency across all servers in the solution. The configuration is designed to achieve a target of 2,400 Exchange Server 2010 IOs for 20,000 users. The results show excellent I/O performance with nearly 2,700 Exchange Server 2010 user I/Os achieved. This provides additional throughput capacity to absorb any unexpected spikes during very heavy user activities.

Table 10. Performance across servers

Database I/O	Value
Target transactional I/Os	2400
Achieved transactional I/Os	2685.223
I/O database reads/sec	1663.382
I/O database writes/sec	1021.837
I/O database average read latency (ms)	17.528
I/O database average write latency (ms)	3.022

Transaction Log I/O	Value
I/O log writes/sec	937.225
I/O log average write latency (ms)	1.179

Database backup and recovery performance

There are two test reports in this category. The first test measures the sequential read rate of the database files, and the second test measures recovery/replay (playing the transaction logs against the databases) performance.

Database read-only performance

This test measures the maximum rate at which databases can be backed up using VSS. The data in Table 11 shows the average rate to back up a single database file across two servers.

Table 11. Database read-only performance

Metric	MBX01	MBX02
MBs read per second per database	32.92	32.84
MBs read per second per server	658.37	656.86
Metric	Value	
Total MBs read per second for all two servers	1315.23	

Transaction log recovery/replay performance

This test measures the maximum rate at which log files can be played against the databases. Table 12 shows the average rate for 500 log files played against a single database. Each log file is 1 MB.

Table 12. Transaction log recovery/replay performance

Metric	Value
Average time to play one log file (in seconds)	4.44

Detailed test results

The detailed test results for both mailbox servers are attached to this PDF in the form of Jetstress HTML reports. These reports include:

- 24-hour stress test
- Checksum for 24-hour stress test
- Two-hour performance test
- Checksum for two-hour performance test
- Database backup test
- Soft recovery test

How to view Jetstress reports

To reveal the list of Jetstress reports (HTML files), click the paper clip icon in the left-hand pane of this document as shown in Figure 4.

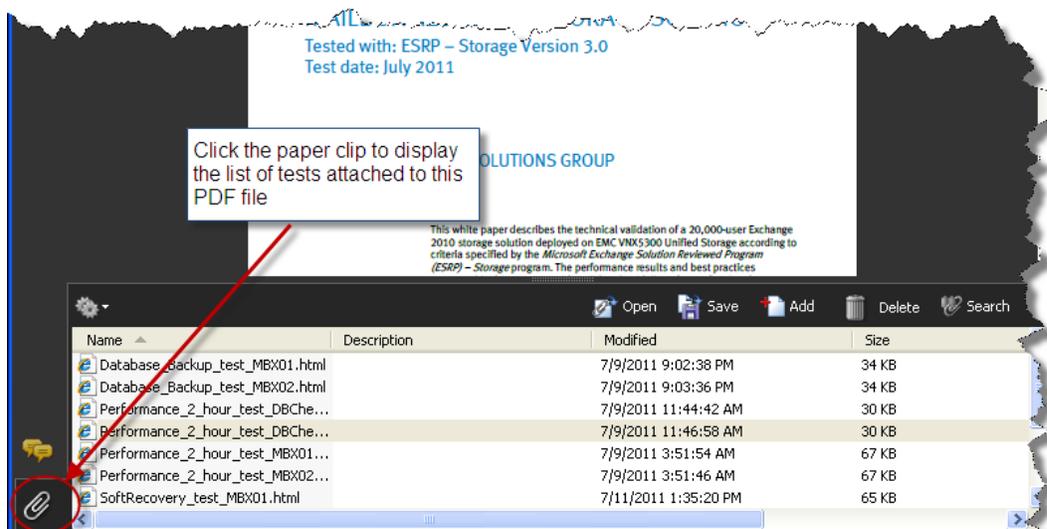


Figure 4. Viewing the Jetstress reports

Double-click on the file you wish view to open the report in your browser, as displayed in Figure 5.

Jetstress 2010 Test Result Report

Microsoft Exchange Jetstress 2010

Test Result Report

Checksum Statistics - All

Database	Seen pages	Bad pages	Correctable pages	Wrong page-number pages	File length / seconds
C:\MP\DB01\Jetstress001001.edb	34097986	0	0	0	1065562 MB/2816
C:\MP\DB02\Jetstress002001.edb	34097474	0	0	0	1065546 MB/2816
C:\MP\DB03\Jetstress003001.edb	34097474	0	0	0	1065546 MB/2832
C:\MP\DB04\Jetstress004001.edb	34097730	0	0	0	1065554 MB/2816
C:\MP\DB05\Jetstress005001.edb	34097730	0	0	0	1065554 MB/2816
C:\MP\DB06\Jetstress006001.edb	34097474	0	0	0	1065546 MB/2820
C:\MP\DB07\Jetstress007001.edb	34097730	0	0	0	1065554 MB/2832
C:\MP\DB08\Jetstress008001.edb	34097474	0	0	0	1065546 MB/2810
C:\MP\DB09\Jetstress009001.edb	34097474	0	0	0	1065546 MB/2840
C:\MP\DB10\Jetstress010001.edb	34097730	0	0	0	1065554 MB/2826
C:\MP\DB11\Jetstress011001.edb	34097730	0	0	0	1065554 MB/2834
C:\MP\DB12\Jetstress012001.edb	34097986	0	0	0	1065562 MB/2816
C:\MP\DB13\Jetstress013001.edb	34097730	0	0	0	1065554 MB/2816
C:\MP\DB14\Jetstress014001.edb	34097474	0	0	0	1065546 MB/2820
C:\MP\DB15\Jetstress015001.edb	34097730	0	0	0	1065554 MB/2834
C:\MP\DB16\Jetstress016001.edb	34097474	0	0	0	1065546 MB/2816
C:\MP\DB17\Jetstress017001.edb	34097730	0	0	0	1065554 MB/2816
C:\MP\DB18\Jetstress018001.edb	34097730	0	0	0	1065554 MB/2820
C:\MP\DB19\Jetstress019001.edb	34097730	0	0	0	1065554 MB/2820

Figure 5. Displaying a report in your browser

Conclusion

The testing and validation of this Exchange Server 2010 mailbox resiliency storage solution demonstrates the following:

- The EMC VNX5300 storage array is an excellent platform for Exchange Server 2010 storage. The array provides optimal performance with adequate capacity for deploying large Exchange mailboxes.
- EMC VNX5300 is fully optimized for Exchange Server 2010 workloads.
- A building block approach simplifies mailbox server design and facilitates scalable, predictable performance for all mailbox servers.

EMC recommends the use of a building block approach when designing storage solutions for Exchange Server 2010. In this solution, a building block of 10,000 users with 100 messages sent/received per user per day, 0.10 IOPS throughput with a 20 percent reserve (0.12 IOPS tested), and a 1.5 GB mailbox size is used.

The configuration is scaled to 20,000 users through the utilization of two building blocks on an EMC VNX5300 storage array. This configuration was shown to meet all of the recommended Microsoft Exchange Server 2010 metrics for performance and capacity. In addition, the performance (IO and latency) test results presented in this Microsoft-approved ESRP document demonstrate that this solution provides significant processing reserve, or “overhead,” to accommodate future user growth.

EMC has published multiple Proven Solutions white papers that demonstrate that the EMC VNX family of storage arrays can handle very heavy Exchange Server 2010 workloads. For more information, visit <http://www.emc.com/exchange>.

Note: This document was developed by EMC and reviewed by the Microsoft Exchange Product team. The test results and data presented in this document are based on the tests introduced in the ESRP test framework. Customers should not quote the data directly for their pre-deployment verification. It is still necessary to go through the exercises to validate the storage design for each specific customer environment. The ESRP program is not designed to be a benchmarking program; tests are not designed to achieve the maximum throughput for a given solution. Rather, the program is focused on producing recommendations from vendors for Exchange applications. The data presented in this document should not be used for making direct comparisons between various solutions.

Contact EMC

EMC recommends that you consult with EMC Professional Services to assist with the design and deployment of a similar solution. For information regarding this or any other EMC solution, use the following numbers:

- United States: (800) 782-4362 (SVC-4EMC)
- Canada: (800) 543-4782 (543-4SVC)
- Worldwide: (508) 497-7901

For additional information on EMC products and services available to customers and partners, visit <http://EMC.com> or <http://Powerlink.EMC.com>.