

EMC VNX5700 UNIFIED STORAGE 60,000 USERS WITH 2 GB MAILBOXES MICROSOFT EXCHANGE SERVER 2010 MAILBOX RESILIENCY STORAGE SOLUTION

Tested with: ESRP – Storage Version 3.0


Test date: June 2011

EMC SOLUTIONS GROUP

Abstract

This paper describes the technical validation of a 60,000 user Exchange 2010 storage solution deployed on EMC VNX5700 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 VNX5700 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 h8813.1

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 VNX5700	6
Solution description	8
Storage design for Exchange Server 2010	9
Building block used in this solution	10
About storage pools	10
Additional factors that can affect mailbox server scalability.....	11
Targeted customer profile	12
Tested deployment	12
Best Practices	15
Mailbox servers.....	15
Networking	16
Core storage.....	16
Backup strategy	17
Information resources	17
Test results summary	17
Reliability (stress)	17
Storage performance.....	18
Individual server metrics	18
Performance across servers.....	18
Database backup and recovery performance.....	19
Database read-only performance.....	19
Transaction log recovery/replay performance	19
Detailed test results	20
How to view Jetstress reports	20
Conclusion	20
Contact EMC	21

Overview

This document provides information about EMC® VNX5700™ Unified Storage array performance with 60,000 Microsoft Exchange Server 2010 users. Performance validation is 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 21.

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

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.

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

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 VNX5700

The EMC VNX5700 model is a member of the VNX series next-generation storage platform, powered by the Intel quad-core Xeon 5600 series processors with a 6 Gb SAS drive back end, providing the industry's highest bandwidth. The EMC VNX5700 model is designed to deliver maximum performance and scalability for enterprises. It is a converged platform that replaces CLARiiON and Celerra and enables organizations to dynamically grow, share, and cost-effectively manage multiprotocol file systems and multiprotocol block storage access.

The VNX Operating Environment allows Microsoft Windows and Linux/UNIX clients to share files in multiprotocol (NFS and CIFS) environments. At the same time, it supports iSCSI, Fibre Channel (FC), and FCoE access for high-bandwidth and latency-sensitive block applications.

Table 1 lists VNX5700 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. VNX5700 features summary (shown for block only)

System feature	Value
Minimum / maximum drives for each system	4 / 500
Drive types	Flash, SAS, NL-SAS
I/O architecture	PCI-e Gen 2 x 4
Disk array enclosure (DAE) options	25 x 2.5" SAS/Flash drives – 2U 15 x 3.5" SAS/Flash drives – 3U
RAID options	0 / 1 / 10 / 5 / 6
Number of storage processors (SPs)	2
Block protocols	FC, iSCSI, FCoE
CPU / memory for each array	Quad-core Intel Xeon 5600 @ 2.40 GHz / 36 GB (18 GB for each SP)
Maximum block flex I/O modules for each array	10
Maximum raw capacity	984 TB
Maximum SAN hosts	4,096
Maximum number of pools	40
Maximum number of LUNs	4,096
Maximum number of ports for each array	24
2/4/8 Gb/s FC maximum ports for each array	24
1GBaseT iSCSI maximum number of ports for each array	16
10GbE iSCSI min / maximum number of ports for each array	12
FCoE maximum number of ports for each array	12
6 Gb/s SAS buses for DAE connections	4
Maximum block Ultraflex™ I/O modules for each array	10
Management and base software	Integrated software includes: <ul style="list-style-type: none"> • Unisphere • Block Compression • Virtual Provisioning™

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 60,000 users in a mailbox resiliency configuration across two intelligent VNX5700 storage arrays. The solution is designed to provide outstanding performance and flexibility for today's and tomorrow's Exchange users.

In this solution, 60,000 users are deployed across two DAGs with six mailbox servers per DAG (30,000 users per DAG). Each DAG has two RAID-protected copies of every Exchange database, a primary (active) copy and a secondary (passive) copy, which are evenly split between the two VNX5700 arrays. Each Exchange database is replicated to an alternate mailbox server on a second VNX5700 array through the use of the native DAG host-based log shipping mechanism.

Each mailbox server in this solution is designed to support up to 10,000 users with a 2 GB mailbox capacity and 0.12 IOPS per user with a 20 percent throughput reserve, also known as "overhead" (0.14 IOPS per user tested).

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. For example, during a mailbox server failure or maintenance operation, each mailbox server is capable of handling the compute and storage requirements of all active databases.

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 six mailbox servers provides service to all 60,000 users.

For more information, access the Windows Hardware Compatibility List link for the EMC VNX5700 unified storage array:

- <http://www.windowsservercatalog.com/item.aspx?itemId=8be6acc5-169c-5b78-03a7-076e05034aa6&bCatID=1282>

The following diagram illustrates the architecture of the EMC VNX5700 60,000 Mailbox Exchange Server 2010 mailbox resiliency storage solution.

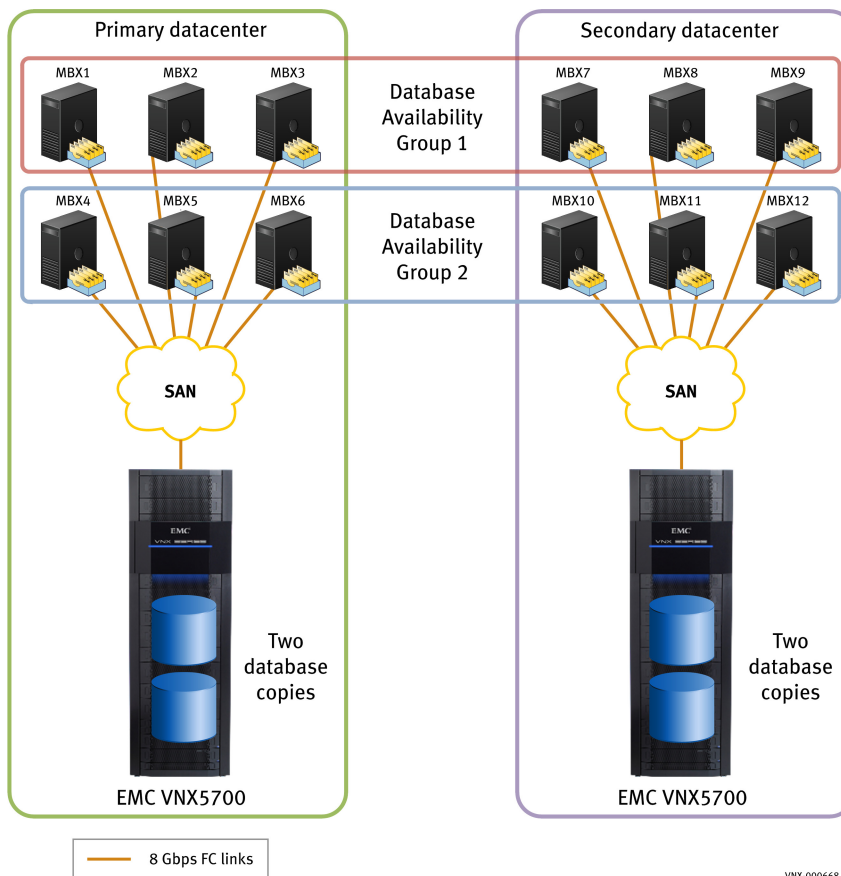


Figure 2. Solution architecture

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 smooth Exchange operation. 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, EMC customers 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 because it makes Exchange environment expansion simple and straightforward. EMC best practices involving the building block approach for Exchange Server design have proven to be very successful in many customer implementations.

Building block used in this solution

The mailbox server building block used in this solution is designed to support 10,000 users per server with a 2 GB mailbox size and 0.12 IOPS per user (0.14 IOPS were tested, which reflects a 20 percent IOPS “overhead,” or reserve).

The building block in this solution is made up of 42 NL SAS drives that support a single mailbox server, which in turn supports 10,000 users. A RAID 1/0 (1+1) RAID group of two 2 TB NL SAS drives houses mailbox server logs. A RAID 1/0 storage pool of 40 2 TB NL SAS drives provides I/O and mailbox capacity for databases. The solution uses six of these building blocks (252 NL SAS drives) on the primary storage array to scale the configuration up to 60,000 users. The six building blocks are duplicated on the secondary storage array.

The following table 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	0.12 IOPS with a 20 percent reserve (0.14 IOPS tested)
Mailbox size	2 GB
Drive type and capacity	2 TB NL SAS drives, 7,200 rpm
RAID type	RAID 1/0
Database LUN size	1.6 TB
Log LUN size	80 GB
Number of disks to support a single mailbox server	42

Storage pools with thick LUNs are used on the VNX5700 array for Exchange databases while traditional RAID groups LUNs are used for logs.

About storage pools

The use of storage pools simplifies storage provisioning. Traditional storage provisioning with only RAID groups restricts the number of disks you can have in a group to 16. Storage pools, on the other hand, enable you to manage potentially hundreds of disks at a time. Such pool-based provisioning provides benefits similar to metaLUN striping across many drives but, unlike metaLUNs, storage pools require minimal planning and management effort. Storage pools support the same RAID protection levels as RAID groups do: RAID 5, RAID 6, and RAID 1/0.

Figure 3 illustrates the building block’s database and log LUN configuration, designed to provide the best performance and sufficient capacity for 10,000 users with a 2 GB mailbox size and 0.12 IOPS per user (0.14 IOPS tested to include overhead). From each RAID 10 storage pool, 20 database LUNs are created. 20 log LUNs are created from each 1/0 (1+1) RAID group.

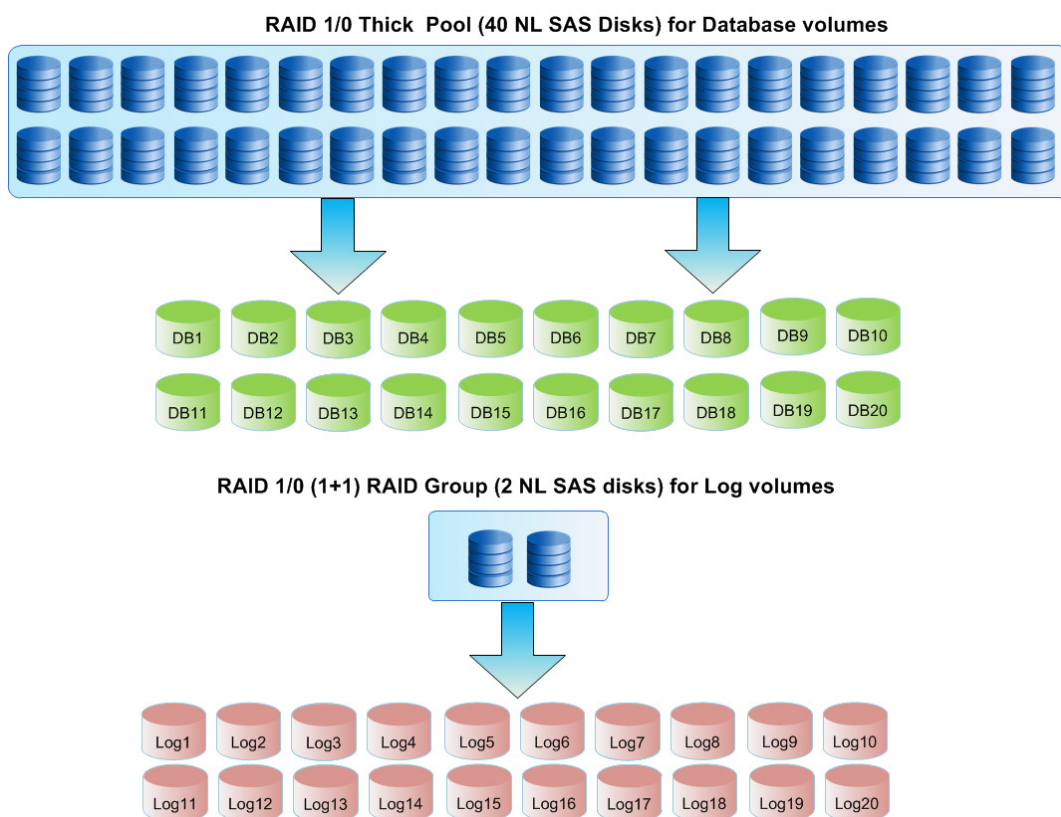


Figure 3. Building block's database and log LUN configuration

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 that must be considered when designing a scalable Exchange solution. Other factors that can affect server scalability are:

- 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.

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>.

For more information on identifying and addressing performance bottlenecks in an Exchange system, refer to Troubleshooting Microsoft Exchange Server Performance, available at <http://technet.microsoft.com/en-us/library/dd335215.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, which supports 60,000 Exchange users, has the following characteristics:

- 60,000 active mailboxes
- 100 messages sent/received per user per day, 0.12 IOPS throughput with a 20 percent reserve (0.14 IOPS tested)
- 2 GB mailbox size
- Six 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 VNX5700 storage arrays (the solution was tested with six mailbox servers connected to a single array)

Note: If more than two 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	60,000
Number of Database Availability Groups (DAGs)	2
Number of servers/DAG	6
Number of active mailboxes/server	10,000
Number of databases/mailbox server	20
Number of copies/database	2

Attribute	Value
Number of mailboxes/database	500
User profile	0.12 IOPS throughput with a 20 percent reserve (0.14 IOPS tested)
Database LUN size	1.6 TB
Log LUN size	80 GB
Total database size for performance testing	150 TB (25 TB for each mailbox server)
Percentage of storage capacity used by Exchange database ³	80 percent (150/192)

Table 4. Storage hardware

Component	Description
Storage Connectivity (FC or iSCSI)	FC
Storage model and OS/firmware revision	VNX5700, VNX Block Operation Environment version: 05.31.000.5.008 http://www.windowsservercatalog.com/item.aspx?itemId=8be6acc5-169c-5b78-03a7-076e05034aa6&bCatID=1282
Storage cache	18 GB mirrored
Number of storage controllers	Two storage processors (SPs)
Number of storage ports	8 (4 for each SP)
Maximum bandwidth of storage connectivity to host	16 Gbps (2 x 8 Gbps HBAs)
Switch type/model/firmware revision	Cisco MDS 9509 Fiber Channel, 8 Gbps, firmware 4.1 (3a)
HBA model and firmware	Qlogic dual-port 8 Gbps HBA QLE2562, firmware version 5.04.01
Number of HBAs/host	Two
Host server type	Dell PowerEdge R900 with Intel(R) Xeon(R) CPU X7350 @ 2.93GHz, 2.92GHz, 4 Core, 128 GB RAM
Number of disks used in solution	252
Maximum number of disks supported by VNX5700 storage array	500

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.

Table 5. Storage software

Component	Description
HBA driver	STOR Miniport 9.1.8.25
HBA QueueTarget Setting	256
HBA QueueDepth Setting	256
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 7,200 RPM
Raw capacity for each disk (GB)	1,834 GB
Number of physical disks in test	240
Total raw storage capacity (GB)	440,160 GB
Raid level	RAID 1/0 EMC VNX storage pool
Total formatted capacity	192 TB (32 TB for each mailbox server)
Storage capacity utilization (percentage)	44 percent (192,000/440,160)
Database capacity utilization (percentage)	34 percent (150,000/440,160)

Table 7. Storage disk configuration (transactional log disks)

Attribute	Description
Disk type, speed, and firmware revision	2 TB NL SAS 7,200 RPM
Raw capacity for each disk (GB)	1,834 GB
Number of physical disks in test	12
Total raw storage capacity (GB)	22,008 GB
Raid level	RAID 1/0
Total formatted capacity	9.6 TB (1.6 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 early versions of Exchange, particularly with regard to IO and storage. There have been many changes to the core schema and the extensible storage engine (ESE) to reduce IO footprint. Due to this IO reduction, Exchange Server 2010 can be deployed on low cost SATA and NL SAS disks in addition to FC, SAS, and SSD drives. Depending on the Exchange Server 2010 deployment model you choose, 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 10 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 recommendations, EMC recommends following 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.
- Install EMC PowerPath® for optimal path management and maximum IO performance. For more information on installing and configuring EMC PowerPath, visit <http://www.emc.com/products/detail/software/powerpath-multipathing.htm>.

- 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.

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.

Core storage

EMC VNX5700 is already optimized for Exchange Server 2010 workloads and requires only a few minor adjustments. Follow these guidelines to ensure optimal storage performance:

- Isolate the Microsoft Exchange server database workload from other IO-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 performance issue.
- Set the VNX storage array page size parameter to 16 KB.
- Allocate 1 GB of available SP memory to read cache, and allocate the remainder of the available memory to write cache.
- Ensure that write cache is enabled for all database and log LUNs.
- Plan for performance even in a failover situation.
- Ensure that IO load and bandwidth are evenly balanced and distributed across VNX5700 front-end ports and back-end buses for failover and load balancing.
- When sizing storage, always prioritize IOPS over capacity. After the IOPS requirements are calculated, assess the capacity that is required.
- After establishing actual required IOPS, apply a 20 percent IO “overhead” factor to build in adequate reserve throughput capability.
- Consider additional throughput requirements imposed by BDM.
- Either RAID groups or storage pools can be used for Exchange data.
- 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.
 - 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.

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 IO and capacity requirements for the target profile. Jetstress simulates Exchange Server 2010 IO 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 IO 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, IOs per second per user) and validate that the disk subsystem is capable of maintaining an acceptable level of performance with such a profile. Both configuration options produce similar results.

Note: The Jetstress tool is designed to test performance of an Exchange storage subsystem before the subsystem is placed in 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 IO 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. 24-hour stress test results reveal that:

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

Storage performance

The purpose of the primary storage performance test is to exercise the storage with the maximum sustainable Exchange IO for two hours with 100 percent user concurrency. The test reveals the amount of time it takes the storage to respond to IO requests under load.

This section presents the sum of all database transactional IOs and the average IO latencies for all databases and logs recorded during the two-hour test duration. Data for each server is presented separately. Aggregates across servers are also included.

Individual server metrics

Table 9 presents the sum of transactional IOs 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 1,400 IOPS for 10,000 users on each server. This includes 20 percent overhead, above the 0.12 IOPS user profile, for reserve throughput capability.

Table 9. Individual server metrics

Database I/O	MBX1	MBX2	MBX3	MBX4	MBX5	MBX6
Achieved transactional I/Os	1450.288	1500.538	1500.414	1547.589	1548.343	1563.900
I/O database reads/sec	906.664	938.185	936.777	965.326	965.691	974.089
I/O database writes/sec	543.626	562.352	563.636	582.264	582.653	589.810
I/O database average read latency (ms)	16.880	16.520	16.538	16.609	16.556	16.522
I/O database average write latency (ms)	12.763	10.143	10.183	11.428	10.295	10.329
Transaction log I/O	MBX1	MBX2	MBX3	MBX4	MBX5	MBX6
I/O log writes/sec	487.937	503.247	498.420	506.042	511.153	513.489
I/O log average write latency (ms)	1.729	1.435	1.439	1.586	1.429	1.435

Performance across servers

Table 10 presents the sum of transactional IOs and the average latency across all servers in the solution. The configuration is designed to achieve a target of 8,400 Exchange Server 2010 IOs for 60,000 users. The results show excellent IO performance with over 9,100 Exchange Server 2010 user IOs 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 IOs	8,400
Achieved transactional IOs	9111.072
IO database reads/sec	5686.732
IO database writes/sec	3424.341
IO database average read latency (ms)	16.604
IO database average write latency (ms)	10.857
Transaction Log I/O	Value
IO log writes/sec	3020.288
IO log average write latency (ms)	1.509

Database backup and recovery performance

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

Database read-only performance

This test measures the maximum rate at which databases can be backed up through the use of VSS. The following data shows the average rate to back up a single database file across six servers.

Table 11. Database read-only performance

Metric	Value
MBs read per second per database	26.85
MBs read per second per server	537.06
Total MBs read per second for all six servers	3,222.36

Transaction log recovery/replay performance

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

Table 12. Transaction log recovery/replay performance

Metric	Value
Average time to play one log file (in seconds)	4.48

Detailed test results

Detailed test results for all six mailbox servers are attached to this PDF in the form of Jetstress HTML reports. Report types are as follows:

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

How to view Jetstress reports

Click the paper clip icon in the left-hand pane of Adobe Reader to reveal the list of Jetstress reports (HTML files). Double-click each file to open the corresponding report in your browser.

Conclusion

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

- The EMC VNX5700 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 VNX5700 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, 0.12 IOPS throughput with a 20 percent reserve (0.14 IOPS tested), and a 2 GB mailbox size is used.

The configuration is scaled to 60,000 users through the utilization of six building blocks on an EMC VNX5700 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>.