

EMC VMAXe 80,000 USERS WITH 2 GB MAILBOXES MICROSOFT EXCHANGE SERVER 2010 MAILBOX RESILIENCY STORAGE SOLUTION


Tested with: ESRP – Storage Version 3.0

Test date: November 2011

Abstract

This white paper describes the technical validation of an 80,000 user Exchange 2010 storage solution deployed on VMAXe 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 VMAXe storage system for a large enterprise Exchange Server 2010 environment.

January 2012



Copyright © 2012 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.

Intel and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Part Number H9524.1

Table of contents

Overview	4
Disclaimer	4
Product Features	4
EMC VMAXe	4
Software suites available	6
Core software	6
Symmetrix VMAXe Advanced Software Suite.....	6
Symmetrix VMAXe Base Software Suite	6
Optional software.....	7
Solution description.....	7
Storage design for Exchange Server 2010	8
Building block used in this solution	9
About Symmetrix Virtual Provisioning	10
Additional factors that can affect mailbox server scalability	11
Targeted customer profile	12
Tested deployment	12
Best Practices	16
Mailbox servers.....	16
Networking	17
Core storage.....	17
Backup strategy	18
Information resources	18
Test results summary	18
Reliability (stress)	19
Storage performance.....	19
Individual server metrics	19
Performance across servers.....	20
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	21
Contact EMC	22

Overview

This document provides information about EMC® Symmetrix® VMAXe™ with Engenuity™ storage array performance with 80,000 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 22.

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.

Product Features

EMC VMAXe

EMC Symmetrix VMAXe, powered by the Intel® Xeon® processor, is a multi-controller, scale-out architecture with consolidation and efficiency for the enterprise. A single VMAXe engine provides the complete foundation for a high-availability storage array. Each VMAXe engine contains two Symmetrix VMAXe directors and redundant interfaces to the EMC Virtual Matrix™ interconnect. Each Symmetrix VMAXe director consolidates front-end, global memory, and back-end functions, enabling direct memory access to data for optimized I/O operations. You can connect up to four VMAXe engines through a set of multiple active fabrics that provide scalable performance and high availability. You can add VMAXe engines non-disruptively to provide linear scale-out of Symmetrix system resources.

The system scales from a single engine (in one system bay) with no additional storage bays to four engines (in four system bays) with three additional storage bays. Figure 1 depicts a fully deployed seven bay VMAXe system with 1080 drives.

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>.

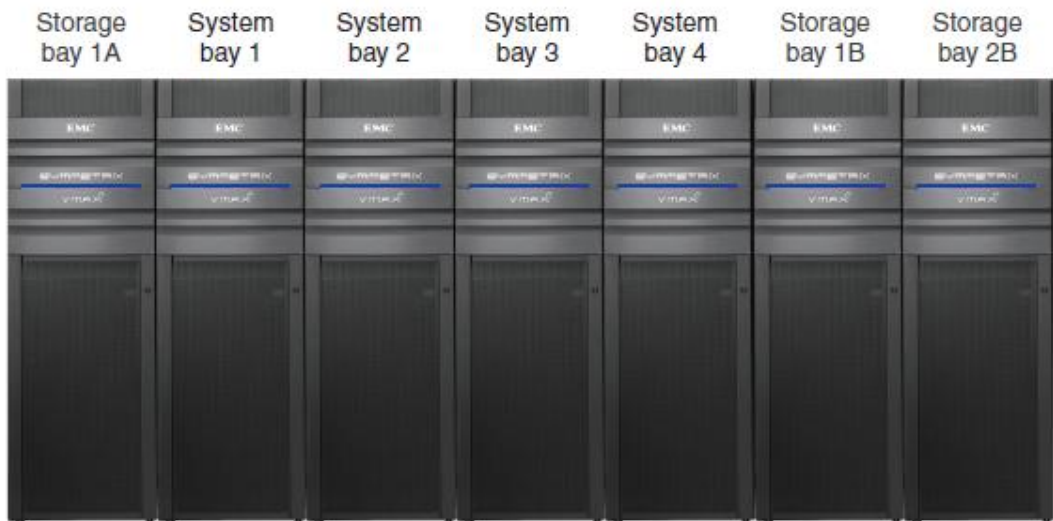


Figure 1. EMC VMAXe

Table 1 lists VMAXe hardware features. For additional VMAXe specifications, visit [www.emc.com](http://www.emc.com/collateral/hardware/specification-sheet/h8716-symmetrix-vmaxe-ss.pdf) at: <http://www.emc.com/collateral/hardware/specification-sheet/h8716-symmetrix-vmaxe-ss.pdf>.

Table 1. VMAXe Hardware features summary

System feature	Value
Minimum / maximum drives for each system	24 / 1080
Drive types	Flash, FC, SATA, SAS
Disk array enclosure (DAE) options	15 x 3.5" drives – 3U
RAID options	0 / 1 / 10 / 5 / 6
Number of VMAXe Engines (director pairs)	1-4
Block protocols	FC, iSCSI, FCoE
CPU per Engine	2 x Quad-core Intel Xeon 2.40 GHz
Maximum Memory per Engine/System	128GB/512GB
Maximum usable capacity per System (4 engines)	1,490 TB
2/4/8 Gb/s FC maximum ports for each Engine/System	16/64
1Gb iSCSI maximum number of ports for each Engine/System	8/32
10Gb iSCSI min / maximum number of ports for each Engine/System	8/32
FCoE maximum number of ports for each Engine/System	8/32

VMAXe delivers uncompromising scalability and flexibility while providing market-leading simplicity and efficiency, and includes the following features:

- Factory pre-configured and built to order to enable fast deployment
- 100% virtually provisioned for ease of use and storage efficiency
- Fully Automated Storage Tiering for Virtual Pools (FAST VP), which optimizes the system for the highest performance and lowest storage cost
- Industry leading replication technologies including SRDFe, TimeFinder, Open Replicator, and an integrated RecoverPoint splitter.
- Simplified management with Symmetrix Management Console (SMC) and Symmetrix Performance Analyzer (SPA)

Intel Xeon 5600 series processors

The EMC VMAXe family utilizes the Intel Xeon 5600 series processors, which help make it 2-3x faster overall than its predecessor does. The VMAXe quad-core processor supports the demands of advanced storage capabilities such as virtual provisioning, compression, and deduplication. Furthermore, performance of the Xeon 5600 series enables EMC to realize its vision for FAST on the VMAXe, with optimized performance and capacity, without tradeoffs, in a fully automated fashion.

Software suites available

VMAXe software includes direct, from-the-factory offerings including core software, an Advanced Software Suite, a Base Software Suite, and optional additional software, making it easier and simpler to attain the maximum overall benefits.

Core software

The VMAXe core software offering includes:

- Enginuity for Symmetrix VMAXe
- Virtual Provisioning
- Symmetrix Management Suite
- Integrated RecoverPoint Splitter
- TimeFinder for VMAXe

Symmetrix VMAXe Advanced Software Suite

The Symmetrix VMAXe Advanced Software Suite option contains all the following software required to run the VMAXe array and includes:

- FAST VP
- Symmetrix Performance Analyzer
- PowerPath SE
- Symmetrix Migration Suite including Open Replicator, Open Migrator, Federated Live Migration

Symmetrix VMAXe Base Software Suite

The Symmetrix VMAXe Base Software Suite contains the core software and all of the Advanced Software Suite components, excluding FAST VP.

Optional software

The Symmetrix VMAXe optional software packages include:

- FAST VP add-on to Base Software Suite
- EMC RecoverPoint
 - Continuous remote replication (CRR)
 - Continuous data protection (CDP)
 - Concurrent local and remote (CLR)
- EMC SRDFe (SRDFe/A, SRDFe/S, SRDFe/DM)
- EMC Replication Manager
- EMC Ionix ControlCenter, and ProSphere Package (Performance Manager, SAN Manager, and StorageScope)
- EMC Ionix Storage Configuration Advisor
- EMC PowerPath
- EMC Open Replicator with Dynamic Mobility

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 80,000 users in a mailbox resiliency configuration across two intelligent two-engine VMAXe storage arrays. The solution is designed to provide outstanding performance and flexibility for today's and tomorrow's Exchange users.

In this solution, 80,000 users are deployed across two DAGs with eight mailbox servers per DAG (40,000 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 VMAXe arrays. Each Exchange database is replicated to an alternate mailbox server on a second VMAXe 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 0.08 IOPS per user (.067 IOPs per user with 20 percent reserved throughput.)

This solution eliminates a single point of failure and handles 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: We can compare this solution architecture to one in which Exchange Server 2010 is deployed in a stand-alone configuration (no DAGs) and a single array with eight mailbox servers provides service to all 80,000 users.

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

<http://www.windowsservercatalog.com/item.aspx?idItem=130a540c-1570-e8a7-87c4-73c6c1db56ed&bCatID=1338>

Figure 2 illustrates the architecture of the VMAXe 80,000 Mailbox Exchange Server 2010 mailbox resiliency storage solution.

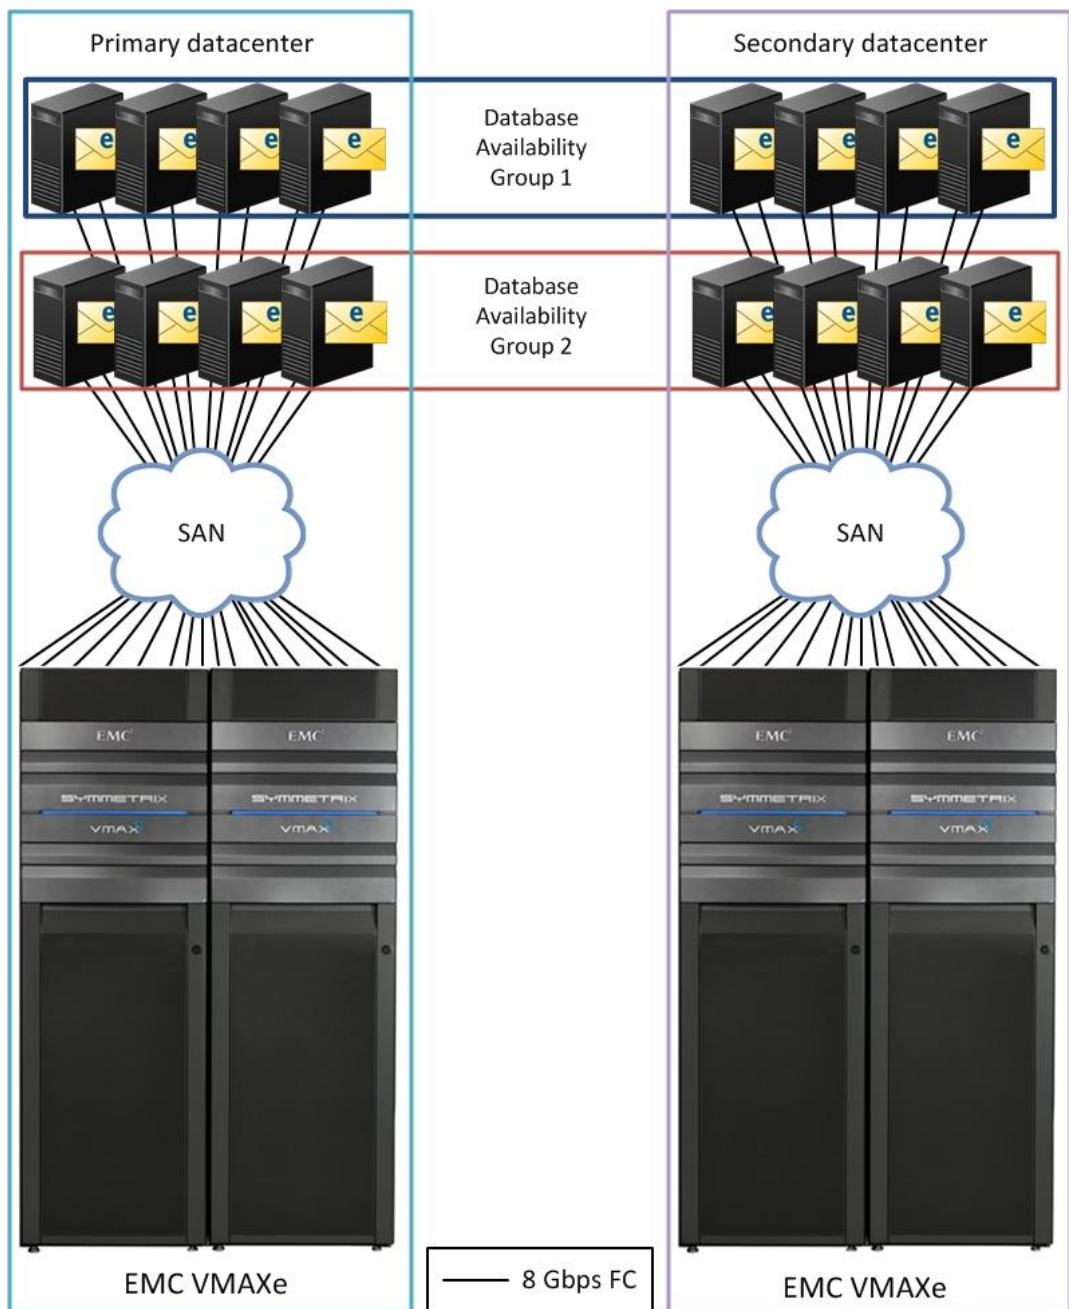


Figure 2. Solution architecture

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 historically used to simplify the sizing and configuration of large amounts of storage 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 0.08 IOPS per user (0.067 IOPs per user with 20 percent reserved throughput.) From a storage perspective each server requires 36 2TB SATA drives in a mirrored configuration to support a single mailbox server.

The Symmetrix VMAXe is a virtually provisioned array, which allows for the pooling of storage resources. Pooling resources enables efficiencies in storage allocations and balances system resources to allow any one server in the environment to achieve optimal performance across all available drives. With these efficiencies in mind we pooled all eight servers into a single common virtually provisioned (thin) pool, where both the log and data storage devices shared the same 300 2 TB SATA drives (6 of which were hot spares.) We configured the pool using RAID1, allocated the storage on demand (as requested from the Exchange servers), and widely striped the storage across the thin pool containing the 300 drives on both the primary and secondary storage arrays.

The EMC validation team configured each server with 10 database LUNs and 10 log LUNs. We configured database LUNs as 2.2 TB 16-way concatenated meta volumes and the log LUNs as 200 GB 4-way concatenated meta volumes.

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	0.08 IOPS (.067 IOPs per user with 20 percent reserved throughput)
Mailbox size	2 GB

Item	Description
Drive type and capacity	2 TB SATA drives, 7,200 rpm
RAID type	RAID 1
Database LUN size	2.2 TB
Log LUN size	200 GB
Number of disks to support a single mailbox server / All mailbox servers	36 / 294 total pooled disks

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 0.08 IOPS per user (0.067 IOPs per user with 20 percent reserved throughput).

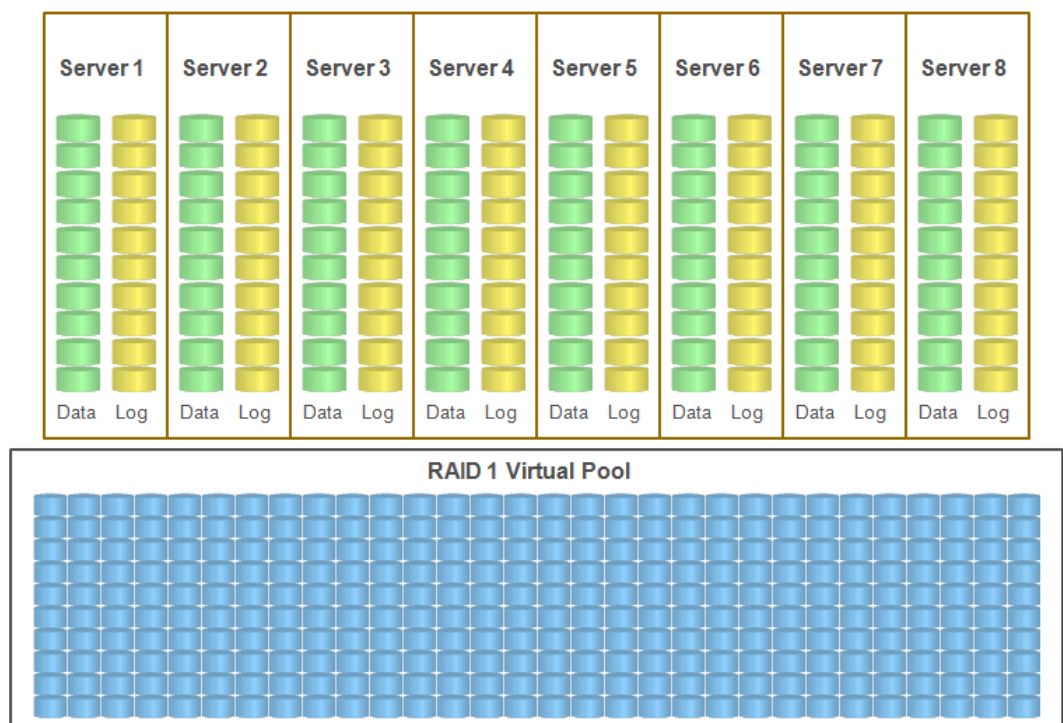


Figure 3. Building block’s database and log LUN configuration

About Symmetrix Virtual Provisioning

Symmetrix Virtual Provisioning introduces a new type of host-accessible device called a “thin” device that you can use in many of the same ways as that of regular host-accessible Symmetrix devices. Unlike regular Symmetrix devices, with thin devices you do not need to completely allocate physical storage at the time you create the device and present that device to a host. The physical storage used to supply drive space for a thin device comes from a shared “thin storage pool” that has been associated with the thin device.

A thin storage pool is comprised of a new type of internal Symmetrix device called a “data device” that is dedicated to providing the actual physical storage used by thin devices. To illustrate the feature in this solution, we configured the data devices as RAID 1 devices. Initially, thin devices are not associated with any particular thin pool.

We must perform an operation referred to as “binding” to associate a thin device with a thin pool. For this solution we bound all data and log thin devices to the same “exchange” pool.

When a write is performed to a portion of the thin device, the Symmetrix array allocates a minimum amount of physical storage from the pool, and maps that storage to a region of the thin device, including the area targeted by the write. These storage allocation operations are performed in small units of storage called “thin device extents.” Thin device extents may also be called “chunks,” “track groups,” or just “extents.” The Symmetrix array uses a round-robin mechanism to balance the allocation of extents across all the enabled data devices in the pool that have unused capacity. The extent size is 12 tracks (768 KB). That means that the initial bind of a thin device to a pool causes one extent, or 12 tracks, to be allocated per thin device. So a four-member thin meta device causes 48 tracks (3078 KB) to be allocated when the meta is bound to a thin pool.

When performing a read on a thin device, the Symmetrix retrieves the data from the appropriate data device in the storage pool to which the thin device is bound. Reads directed to an unmapped area of a thin device do not trigger allocation operations. The result of reading an unmapped block is that a block in which each byte is equal to zero will be returned. When more storage is required to service existing or future thin devices, you can add data devices to existing thin storage pools. You can also create new thin devices and associate them with existing thin pools.

As noted, it is possible to present a thin device to a host before mapping all of the device’s reported capacity, which is what we did when testing the solution described in this document. It is also possible for the sum of the reported capacities of the thin devices using a given pool to exceed the available storage capacity of the pool. This is called an oversubscribed thin device 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 the ESRP—Storage 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 valid under the Windows Server Virtualization Validation Program. If you deploy Exchange mailbox servers as virtual machines (VMs), the number of users 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 80,000 Exchange users, has the following characteristics:

- Up to 80,000 active mailboxes
- 80 messages sent/received per user per day, 0.08 IOPS throughput (.067 IOPs per user with 20 percent reserved throughput)
- Two GB mailbox size
- Eight 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 VMAXe storage arrays (the solution was tested with eight mailbox servers connected to a single array)

Note: If you require more than two database copies, 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	80,000
Number of Database Availability Groups (DAGs)	2
Number of servers/DAG	8
Number of active mailboxes/server	10,000
Number of databases/mailbox server	10

Attribute	Value
Number of copies/database	2
Number of mailboxes/database	1,000
User profile	0.08 IOPS (.067 IOPs per user with 20 percent reserved throughput.)
Database LUN size	2.2 TB
Log LUN size	200 GB
Total database size for performance testing	192 TB (24 TB for each mailbox server)
Percentage of storage capacity used by Exchange database ²	75 percent (192/257)

Note: The database LUN size tested allowed for 10% reserved capacity for the database file system. The thin pool had capacity available to allow for 20% or more reserved file system capacity to meet Microsoft guidance or unique customer requirements.

Table 4. Storage hardware

Component	Description
Storage Connectivity (FC or iSCSI)	FC
Storage model and OS/firmware revision	VMAXe, Enginuity 5875.231.172
Storage cache	86 GB (usable) mirrored
Number of storage controllers	Two VMAXe engines = 4 controllers
Number of storage ports	32 (16 used for testing)
Maximum bandwidth of storage connectivity to host	16 Gbps (2 x 8 Gbps HBAs)
Switch type/model/firmware revision	Brocade DCX Fiber Channel, 8 Gbps, firmware v6.1.0j
HBA model and firmware	Emulex LPe12000-E 2.00a4
Number of HBAs/host	Two
Host server type	Dell PowerEdge R610 with Intel(R) Xeon(R) CPU E5502 @ 1.86GHz, 1.86GHz, 2 Core, 32 GB RAM
Number of disks used in solution	300
Maximum number of disks supported by VMAXe storage array	1080

² 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 2.41.003
HBA QueueTarget Setting	32
HBA QueueDepth Setting	32
Multi-pathing	EMC PowerPath 5.5 SP1
Host OS	Microsoft Windows Server 2008 R2 Enterprise SP1
ESE.dll file version	14.01.0218.012
Replication solution name/version	N/A

Table 6. Storage disk configuration (mailbox store and transaction log disks)

Attribute	Description
Disk type, speed and firmware revision	2 TB SATA 7,200 RPM
Raw capacity for each disk (GB)	1,882 GB
Number of physical disks in test	300 (6 hot spares)
Total raw storage capacity (GB)	553,308 GB
Raid level	RAID 1 EMC thin pool
Total formatted capacity	257 TB (32 TB for each mailbox server)
Storage capacity utilization (percentage)	46 percent (257,000/553,308)
Database capacity utilization (percentage)	32 percent (176,000/553,308)
Log capacity utilization (percentage)	3 percent (16,000/553,308)

Table 7. 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 I/O and storage. There have been many changes to the core schema and the extensible storage engine (ESE) to reduce the I/O footprint. Due to this I/O reduction, you can 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 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. Exchange Server 2010 supports all RAID types (RAID 0, RAID 5, RAID 10, and RAID 6) 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 VMAXe series storage performance with Exchange Server 2010.

Mailbox servers

To ensure the best possible mailbox server performance, follow these recommendations:

- 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 I/O 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, We recommend that you use multiple physical NICs connected to different networks to isolate user MAPI traffic from database replication traffic.

For iSCSI deployments, use multiple network switches for fault tolerance and performance. Where this is not possible, use VLANs to isolate the iSCSI traffic.

Core storage

EMC VMAXe is already optimized for Exchange Server 2010 workloads. Use the following guidelines to ensure optimal storage performance:

- Always calculate spindle sizing for IOPS requirements before sizing for capacity. I/O requirements include user operations (send/receive) and other overhead (such as BlackBerry) plus an additional 20 percent (at a minimum) to account for log writes, DAG replication, and BDM I/O.
- Where you cannot use large database sizes (over 1 TB), calculate Exchange Server 2010 BDM disk throughput operations, when enabled, separately as a part of your sizing estimates. For VMAXe, sizing should include a 7.68 MB/s (30 IOPs @ 256 KB) overhead factor for each database.
- For Microsoft Exchange Server 2010, we recommend that you share database and log volumes across the same VMAXe spindles.
- While sharing spindles is acceptable for databases and logs on the back end of the VMAXe, we recommend that you separate database and log onto their own separate LUNs and file systems on the hosts.
- Evenly distribute Exchange data across the VMAXe back end and disk resources. Virtual provisioning, native to the VMAXe, helps to ensure a balanced configuration across all resources. When adding drives to a pool, you can use automated pool rebalancing to redistribute data evenly across the system.
- Base the number of thin pools used in a configuration on high availability or performance isolation requirements. For example if you maintain multiple database copies on the same VMAXe, you should maintain those database copies on their own thin pool. If isolation is not a requirement, we recommend that you use large thin pools, where multiple servers of like workload share the same spindles. You can use technologies like FAST/VP to help optimize performance for configurations where data is shared across common pools and multiple tiers of storage are available.
- Balance front end processor and port utilization across all available VMAXe resources intended for the Exchange environment. For Fibre Channel, use port 0 of a given front end processor (a slice) before also using port 1 of the same processor.

- You can use concatenated meta volumes can be used in conjunction with Exchange Server 2010 for both the database and log volumes. You can easily expand concatenated meta volumes non-disruptively when you require additional host addressable space.
- You can also use striped meta volumes for Exchange Server 2010, which will provide the best possible performance. You can also expand striped meta volumes non-disruptively; however, you will need additional space in order to perform the expansion.

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.

We validated this solution 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).

You can configure Jetstress to:

- Test the maximum I/O throughput available to the disk subsystem within the performance constraints of Exchange.
- Accept a specific user profile (user count, I/Os 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 I/O 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 a high I/O load for a long period of time. The stress test analyzes both log and database files for integrity after the stress tests to ensure that there is no database or log corruption. Our 24-hour stress test results reported:

- No errors in the saved event log file.
- No errors 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 I/O for two hours with 100 percent user concurrency. The test reveals the amount of time it takes the storage to respond to I/O requests under load.

This section presents the sum of all database transactional I/Os and the average I/O latencies for all databases and logs recorded during the two-hour test duration. The tables below present data for each server separately and include aggregates across servers.

Individual server metrics

Table 8 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 800 IOPS for 10,000 users on each server.

Table 8. Individual server metrics

Database I/O	MBX1	MBX2	MBX3	MBX4	MBX5	MBX6	MBX7	MBX8
Achieved transactional I/Os	819.0	815.6	819.3	817.5	823.2	828.8	806.1	811.0
I/O database reads/sec	514.4	512.5	514.0	513.1	516.8	519.9	506.5	509.0
I/O database writes/sec	304.6	303.1	305.3	304.4	306.4	308.8	299.6	302.0
I/O database average read latency (ms)	19.2	19.1	19.3	18.9	19.0	19.0	19.2	19.2
I/O database average write latency (ms)	8.6	8.2	8.7	9.5	9.4	9.6	8.2	8.4
Transaction log I/O	MBX1	MBX2	MBX3	MBX4	MBX5	MBX6	MBX7	MBX8
I/O log writes/sec	258.7	257.6	259.1	258.7	261.7	264.0	257.5	259.0
I/O log average write latency (ms)	0.76	0.76	0.79	0.71	0.64	0.65	0.79	0.77

Performance across servers

Table 9 presents the sum of transactional I/O and the average latency across all servers in the solution. The configuration is designed to achieve a target of 6,400 Exchange Server 2010 I/Os for 80,000 users. The results show excellent I/O performance that achieve over 6,500 Exchange Server 2010 user I/Os.

Table 9. Performance across servers

Database I/O	Value
Target transactional I/Os	6,400
Achieved transactional I/Os	6,541
I/O database reads/sec	4,107
I/O database writes/sec	2,434
I/O database average read latency (ms)	19.1
I/O database average write latency (ms)	8.8
Transaction Log I/O	Value
I/O log writes/sec	2,076
I/O log average write latency (ms)	0.73

Database backup and recovery performance

There are two tests reported in this category. The first test measures the sequential read rate of database files, and the second test measures recovery/replay (playing 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 eight servers.

Table 10. Database read-only performance

Metric	Value
MBs read per second per database	23.96
MBs read per second per server	239.6
Total MBs read per second for all eight servers	1,917

Transaction log recovery/replay performance

This test measures the maximum rate at which log files can be played against databases. Table 11 shows the average rate for 500 log files played against a single database. Each log file is 1 MB in size.

Table 11. Transaction log recovery/replay performance

Metric	Value
Average time to play one log file (in seconds)	3.15

Detailed test results

Detailed test results for all eight mailbox servers are attached to this PDF in the form of Jetstress HTML reports, and include the following:

- 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) attached to this document. 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 VMAXe storage array powered by the Intel Xeon processor is an excellent platform for Exchange Server 2010 storage. The array provides optimal performance with adequate capacity for deploying large Exchange mailboxes.
- EMC VMAXe 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.
- Large storage pools help to simplify administration and deployment as well as optimize storage utilization across all available resources.

EMC recommends the use of a building block approach when designing storage solutions for Exchange Server 2010. This solution uses a building block of 10,000 users with 80 messages sent/received per user per day, 0.08 IOPS throughput (.067 IOPs per user with 20 percent reserved throughput,) and a 2 GB mailbox size.

The configuration is scaled to 80,000 users through the utilization of eight building blocks on an EMC VMAXe storage array. This configuration meets all of the recommended Microsoft Exchange Server 2010 metrics for performance and capacity.

EMC has published multiple Proven Solutions white papers that demonstrate that EMC 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 of validating 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. Do not use the data presented in this document 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 the website <http://EMC.com> or <http://Powerlink.EMC.com>.