

EMC SYMMETRIX VMAX PERFORMANCE REVIEW FOR MICROSOFT EXCHANGE SERVER 2013

With EMC Virtual Storage Integrator (VSI) and VMware vSphere

- Storage building-block design
- Automated performance optimization with FAST VP
- Simplified storage provisioning and management with VSI

EMC Solutions

Abstract

This white paper validates the performance of a large-scale, virtualized Microsoft Exchange 2013 configuration with EMC® Symmetrix® VMAX® 20K and VMware vSphere 5.1. This paper provides best practices and guidelines that simplify Exchange 2013 storage design by using a building block approach. This solution also includes EMC Unisphere® for VMAX and EMC Virtual Storage Integrator (VSI) for VMware vSphere for simplified storage management and performance monitoring.

November 2013



Copyright © 2013 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 H12390

Table of contents

Executive summary	5
Business case	5
Solution overview	5
Key results	5
Introduction	6
Purpose	6
Scope	6
Audience	6
Terminology	6
Technology Overview	8
Overview	8
EMC Symmetrix VMAX series	8
VMAX hardware support	8
VMAX features	8
EMC Unisphere for VMAX	9
EMC Virtual Storage Integrator for VMware vSphere	10
EMC PowerPath/VE	11
Microsoft Exchange Server 2013	11
VMware ESXi	11
Solution architecture and configuration	12
Overview	12
Solution architecture	12
Targeted Exchange user profile	13
Hardware	13
Software	14
Exchange 2013 storage design and configuration on VMAX	15
Overview	15
Building-block design method	15
Building block design details in this solution	16
Phase 1: Collect user requirements	16
Phase 2: Design the storage architecture based on user requirements	16
Phase 3: Validate the design	19
Exchange storage design best practices on VMAX	19
Front-end port usage and zoning	20
Thin pool design and configuration	21
Symmetrix Virtual Provisioning	21

Thin pool design.....	22
FAST VP design and configuration	22
Symmetrix FAST VP	22
FAST VP design.....	23
Using Unisphere to configure FAST VP	23
Using VSI for storage provisioning.....	25
Solution Validation and Test Results	28
Overview.....	28
Test scenarios.....	29
Baseline test result	29
Jetstress result	29
Using Unisphere to analyze performance.....	30
FAST VP test result	32
Jetstress result	32
Performance improvement between Jetstress baseline and FAST VP tests	33
Using Unisphere to check FAST VP status	34
Conclusion	35
Summary	35
Findings.....	35
References.....	36
White papers	36
Product documentation.....	36
Other documentation	36

Executive summary

Business case

Today email plays a critical role in business communications. Microsoft Exchange administrators are challenged with maintaining the highest possible levels of performance and application efficiency. At the same time, organizations are struggling to keep pace with relentless levels of data growth while working with diminishing or stagnant budgets.

EMC has joined forces with industry leading providers of IT infrastructure to create a complete virtualization solution that accelerates deployment of Microsoft Exchange.

This white paper validates the use of EMC Symmetrix VMAX to exceed performance and scalability objectives for a Microsoft Exchange deployment. This solution uses EMC Unisphere for VMAX and EMC Virtual Storage Integrator (VSI) for VMware vSphere to simplify storage management and accelerate implementation.

Solution overview

This document presents a building-block design for a large enterprise Microsoft Exchange 2013 storage solution that includes VMware vSphere 5.1 as the virtualization platform and EMC Symmetrix VMAX 20K as the back-end storage array.

EMC Symmetrix Virtual Provisioning alleviates overprovisioning, which is one of the biggest challenges facing Exchange storage administrators. Virtual Provisioning allocates physical storage only when storage writes occur. This gives administrators more flexibility to estimate future growth and reduces the initial cost of provisioning storage dedicated to Exchange.

EMC Symmetrix Fully Automated Storage Tiering for Virtual Pools (FAST™ VP) automates tiered storage strategies in virtual provisioning environments by easily moving workloads between Symmetrix storage tiers as changes occur in performance over time. FAST VP provides intelligent and efficient storage management for Microsoft Exchange.

Furthermore, the solution eases provisioning and simplifies management of the infrastructure environment with EMC Unisphere for VMAX and Virtual Storage Integrator (VSI) for VMware vSphere.

Key results

This EMC tested solution produces the following key results:

- Virtualized Exchange 2013 building-block design for the test environment is easy to adopt and scales well.
- Sustainable high performance with VMAX 20K and Exchange 2013 was validated using Microsoft Exchange Server Jetstress.
- Virtual Provisioning minimizes initial costs of storage deployment for Exchange mailboxes and room for future growth.
- FAST VP provides intelligent data movements and significantly improves performance.

Introduction

Purpose

This white paper validates the performance of a large-scale, virtualized Microsoft Exchange 2013 configuration with EMC® Symmetrix® VMAX® 20K and VMware vSphere 5.1. It provides best practices and guidelines that simplify Exchange 2013 storage design by using a building block approach.

Scope

This document provides and demonstrates the following:

- Design guidelines and best practices for configuring Exchange 2013 storage on EMC Symmetrix VMAX 20K
- Tested and validated storage building-block design
- VMAX 20K performance validation with Exchange 2013 type input/output (I/O) by using the Microsoft Jetstress tool
- Increased productivity and simplified storage management and provisioning with Unisphere® for VMAX and EMC VSI for VMware vSphere

Audience

This white paper is intended for EMC employees, partners, and customers, including IT planners, storage architects, Exchange administrators, and EMC field personnel that deploy this solution for customers. This paper assumes that you are familiar with all of the components included in this solution.

Terminology

This white paper includes the following terminology.

Table 1. Terminology

Term	Definition
AD	Active Directory
BDM	Background Database Maintenance
CAS	Client Access Server
DAG	Database Availability Group
DNS	Domain Name System
FAST VP	Fully Automated Storage Tiering for Virtual Pools
FQDN	Fully Qualified Domain Name
IOPS	Input/Output operations per second
NFS	Network File System
NIC	Network interface card
NLB	Microsoft Network Load Balancing
NL-SAS	Near-line serial-attached SCSI
NMP	VMware Native Multipath Plug-in
NTFS	New Technology File System

Term	Definition
ReFS	Resilient File System
rpm	Revolutions per minute
SAS	Serial-attached SCSI
SATA	Serial ATA (Advanced Technology Attachment)
SNMP	Simple Network Management Protocol
SRDF	Symmetrix Remote Data Facility
VMDK	VMware Virtual Machine Disk
VMFS	VMware Virtual Machine File System

Technology Overview

Overview

This section provides an overview of the following key technologies that this solution uses:

- EMC Symmetrix VMAX series
- EMC Unisphere for VMAX
- EMC Virtual Storage Integrator (VSI) for VMware
- EMC PowerPath/VE
- Microsoft Exchange 2013
- VMware ESXi 5.1

EMC Symmetrix VMAX series

The EMC Symmetrix VMAX series provides high-end storage for a data center. The system can scale up to a massive 4 PB and can consolidate more workloads with a much smaller footprint. EMC Symmetrix Virtual Matrix Architecture seamlessly scales performance, capacity, and connectivity on demand to meet all application requirements.

VMAX hardware support

Table 2 lists the tested VMAX systems and the required hardware specifications for these systems.

Table 2. VMAX hardware specifications

	VMAX 10K	VMAX 20K	VMAX 40K
Maximum Disk Drives	1,560	3,200	
Usable Capacity	1.5 PB	2 PB	4 PB (3.5" drives) and 2.8 PB (2.5" drives)
Connectivity	4Gb FC, 8Gb FC, 1GbE, 10GbE, 10Gb FCoE, iSCSI	4Gb FC, 8Gb FC, FICON, 1GbE, 10GbE, 10Gb FCoE, iSCSI	
Front-end Ports	64	128	
Local Replication	EMC TimeFinder, EMC RecoverPoint		
Remote Replication	Symmetrix Remote Data Facility (SRDF), RecoverPoint		

VMAX features

Table 3 describes the features and benefits that the VMAX series provides.

Table 3. EMC Symmetrix VMAX series features

Feature	Benefit
Symmetrix FAST and FAST VP	Automates storage tiering to lower costs and deliver higher service levels
Linear scale-out of storage resources	Consolidates multiple arrays into a single Symmetrix VMAX system

Feature	Benefit
Up to 4 PB usable capacity	Seamlessly scales from 48 to 3,200 drives
One to eight VMAX engine scaling	Consolidates more workloads in a smaller footprint with up to eight highly available Symmetrix VMAX engines
Virtual logical unit number (LUN) technology	Transparently moves data to the correct tiers and RAID types as needed
Virtual provisioning	Simply and efficiently allocates, increases, and reclaims storage
Extended distance protection	Replicates data over extended distances and provides zero-data-loss protection
Information-centric security	Provides built in advanced RSA security technology to keep your data safe, reduce risk, and improve compliance

The EMC VMAX web page at <http://www.emc.com/storage/symmetrix-vmax/symmetrix-vmax.htm> has additional details about VMAX systems.

EMC Unisphere for VMAX

EMC Unisphere for VMAX is an advanced graphical user interface (GUI) for managing Symmetrix VMAX arrays. Unisphere for VMAX enables you to provision, manage, and monitor Symmetrix VMAX systems from one screen and significantly reduce storage administration time.

You can use Unisphere for VMAX to do the following:

- Manage user accounts and roles
- Perform configuration operations (create volumes, mask volumes, set Symmetrix attributes, set volume attributes, set port flags, and create SAVE volume pools)
- Manage volumes (change volume configuration, set volume status, and create or dissolve meta volumes)
- Perform and monitor replication operations
- Manage advanced Symmetrix features, such as FAST VP, enhanced Virtual LUN technology, and so on
- Monitor alerts, including the ability to configure external alert notifications such as email, Simple Network Management Protocol (SNMP), and system logs

With the performance monitoring option, Unisphere for VMAX provides tools for analysis and historical trending of Symmetrix system performance data. Figure 1 shows the Unisphere for VMAX user interface.




Figure 1. Unisphere for VMAX user interface

EMC Virtual Storage Integrator for VMware vSphere

EMC Virtual Storage Integrator (VSI) for VMware vSphere is a plug-in for the vSphere client that provides a single management interface for managing EMC storage within the vSphere environment. You can add and remove features from VSI independently, which provides flexibility for customizing VSI user environments. Features are managed by using the VSI Feature Manager. VSI provides a unified user experience, which enables new features to be introduced rapidly in response to changing customer requirements.

We¹ used the following features during validation testing:

- **Storage Viewer:** Extends the functionality of the vSphere Client for the discovery and identification of VMAX storage devices on vSphere hosts and virtual machines. Storage Viewer presents the underlying storage details to the virtual datacenter administrator, merging the data of several different storage mapping tools into a few seamless vSphere Client views.
- **Unified Storage Management:** Simplifies storage administration of VMAX storage systems. It enables VMware administrators to provision new network file system (NFS) and virtual machine file system (VMFS) datastores, and raw device mapping (RDM) volumes seamlessly within vSphere client.

The *EMC VSI for VMware vSphere Product Guide* on [EMC Online Support](#) has more information.

¹ In this paper, “we” represents the EMC Solutions Group that tested and validated this solution.

EMC PowerPath/VE EMC recommends installing PowerPath/VE on VMware ESXi hosts for advanced multipath functionality, such as intelligent path testing and performance optimization.

EMC PowerPath/VE provides intelligent, high-performance path management with path failover and load balancing optimized for EMC and selected third-party storage systems. PowerPath/VE supports multiple paths between a vSphere host and an external storage device. Having multiple paths enables the vSphere host to access a storage device, even if a specific path is unavailable. Multiple paths can also share the I/O traffic to a storage device. PowerPath/VE is particularly beneficial in highly available environments because it can prevent operational interruptions and downtime. The PowerPath/VE path failover capability avoids host failure by maintaining uninterrupted application support on the host if a path failure occurs (if another path is available).

**Microsoft
Exchange Server
2013**

Microsoft Exchange Server 2013 is an enterprise email and communication system that enables businesses and customers to collaborate and share information. EMC enhances Exchange Server 2013 with a selection of storage platforms, software, and services.

Exchange 2013 builds upon the Exchange 2010 architecture and was redesigned to simplify scaling, hardware utilization, and failure isolation. Exchange 2013 uses Database Availability Groups (DAGs) and mailbox database copies, along with other features, such as single item recovery, retention policies, and lagged database copies, to provide high availability site resilience and Exchange native data protection. The high availability platform, the Exchange Information Store, and the Extensible Storage Engine (ESE) have all been enhanced to provide greater availability, easier management, and reduced costs.

VMware ESXi

With VMware vSphere 5.1, you can virtualize physical computer resources, such as the CPU, RAM, hard disks, and network controllers. This transformation creates fully functional virtual machines that run isolated and encapsulated operating systems and applications the same as physical computers.

The high-availability features of vSphere 5.1, such as vMotion and Storage vMotion, enable seamless migration of virtual machines and stored files from one vSphere server to another with minimal or no performance impact. Coupled with vSphere Distributed Resource Scheduler (DRS) and Storage DRS, virtual machines have access to the appropriate resources at any point in time through load balancing of compute and storage resources.

The VMware Native Multipathing Plug-in (NMP) is the module that vSphere uses by default for multipathing. It provides a default path selection algorithm based on the array type. NMP associates physical paths with a specific storage device or logical unit number (LUN). The specific details for handling path failover for a given storage array are delegated to a Storage Array Type Plug-In (SATP). The specific details for determining which physical path is used to issue an I/O request to a storage device are handled by a Path Selection Plug-In (PSP). SATPs and PSPs are sub plug-ins within the NMP module.

Solution architecture and configuration

Overview

This solution validates an Exchange 2013 infrastructure for customers that need to deploy up to 100,000 mailboxes. This chapter provides an overview of the solution architecture, the targeted Exchange user profile, and the hardware and software used in this solution.

Solution architecture

Figure 2 depicts the solution architecture of the simulated test environment. The solution design represents an Exchange 2013 environment in a mailbox resiliency configuration across two intelligent EMC VMAX 20K storage arrays. You can also deploy all the Exchange mailboxes and databases on a single VMAX array.

In this solution, we deployed 100,000 user mailboxes across two Database Availability Groups (DAGs) and setup all Exchange Mailbox servers as virtual machines in a VMware vSphere 5.1 virtualization environment. Each Exchange mailbox database has two DAG copies, a primary (active) copy and a secondary (passive) copy, with the active copy replicated to the passive one on an alternate Exchange Mailbox server, through the use of the native DAG host-based log shipping mechanism.




Figure 2. Simulated solution architecture

Targeted Exchange user profile

This solution is designed for an enterprise that is planning to deploy its Exchange 2013 environment into a highly reliable and scalable storage system.

We used the Microsoft Jetstress tool to test the Exchange performance by simulating a failover scenario where all 100,000 users were active on ten Mailbox servers and one VMAX storage array in the primary datacenter. This is to verify that the VMAX array can support Exchange performance requirements even in a failover situation.

Table 4 describes the targeted Exchange user profile in this solution. [Exchange 2013 storage design and configuration on VMAX](#) on page 15 has more details about the configuration, and [Solution Validation and Test Results](#) on page 28 has more information on how we validated the storage design.

Table 4. Simulated customer profile used in the test environment

Item	Value
Number of Exchange 2013 mailboxes	100,000
Mailbox IOPS profile (in DAG configuration)	0.1 IOPS per mailbox
Mailbox size	Start at 1 GB, grow to 2 GB
Number of Mailbox Server virtual machines	20 (10 tested)
Number of hypervisor hosts	10 (5 tested)
Number of DAGs and database copies	2 DAG with 2 copies
Number of users per Mailbox Server	10,000 total mailboxes (5,000 active and 5,000 passive during normal operating conditions)
Target average message size	75 KB
BDM configuration	Enabled as 24 x 7

Hardware

Table 5 details the hardware resources used in this solution.

Table 5. Hardware resources

Equipment	Configuration
Storage system	VMAX 20K storage array 4 engines 256 GB system cache (mirrored)
SAN switches	8 Gb/s Fibre Channel switches
Physical server	Rack servers: - CPU: E7-2870, 2.4 GHz, two sockets and 10 cores per socket - Memory: 160 GB

Software

Table 6 details the software resources used in this solution.

Table 6. Software resources

Resource	Version	Purpose
EMC VMAX 20K operating environment	5876.229.145	VMAX operating environment
VSI for VMware vSphere	5.5	New storage provisioning for ESXi hosts and enhanced storage views
EMC PowerPath/VE	5.8	Advanced multipathing for ESXi host
Windows Server	2012	Operating system for Exchange 2013
VMware ESXi	5.1	Hypervisor software

Exchange 2013 storage design and configuration on VMAX

Overview

From a storage architecture perspective, Exchange 2013 offers significant changes and improvements for storing service, database schema, and search operations. These changes allow for more storage design options and the use of larger-capacity, slow-rotating disk drives for Exchange data.

Sizing and configuring storage for Exchange 2013 is a complicated process driven by many variables and requirements that vary between organizations. Properly configured Exchange storage, combined with optimally sized server and network infrastructures, can guarantee smooth Exchange operation and an excellent user experience. One of the methods that can be used to simplify the sizing of large amounts of storage on EMC Symmetrix VMAX series storage arrays for use with Exchange is to define a unit of measure—a mailbox server building block.

This section provides design guidance for creating a virtualized Exchange 2013 building-block. It also describes the front-end port zoning, thin pool and FAST VP design and configuration used in this solution.

Building-block design method

A mailbox server building block represents the amount of storage (I/O, capacity, and bandwidth), server (CPU, memory), and network resources required to support a specific number of Exchange users. The amount of required resources is derived from specific user-profile types, mailbox sizes, and disk requirements. Using the building block approach simplifies the design and implementation of Exchange.

Once the initial building block is designed, it can be reproduced to support the required number of users in your enterprise. By using this approach, EMC customers can create their own building blocks based on their company's specific Exchange environment requirements. This approach is 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 design have proven to be very successful in many customer implementations.

Designing a building block that is appropriate for your specific environment include the following three phases:

- Phase 1: Collect user requirements
- Phase 2: Design the storage architecture based on user requirements
- Phase 3: Validate the design

The following section presents the building block design process for this solution.

Phase 1: Collect user requirements

The solution design represents an Exchange 2013 environment that supports 100,000 users in a mailbox resilient configuration where each mailbox database has two DAG copies for high availability. The user mailboxes have a 2 GB capacity (1 GB as the initial size) and 0.1 IOPS profile (150 messages sent/received per user per day). Table 4 details the Exchange user requirements.

Phase 2: Design the storage architecture based on user requirements

Based on the user requirements identified in Phase 1, you can design an Exchange 2013 configuration for your test environment. In this configuration, VMAX 20K protects and replicates each Exchange database with a different Exchange Mailbox Server member of the same Exchange DAG.

The 100,000 user mailboxes in our solution are added to two DAGs and evenly distributed across twenty Mailbox servers, as shown in Figure 2. These twenty Mailbox servers are configured as virtual machines across ten ESXi servers with two virtual machines on each hypervisor host. Each Mailbox server virtual machine is designed to accommodate a failure or maintenance tasks of other Exchange Mailbox server virtual machines. During a failure or maintenance, each Mailbox server virtual machine is sized to accommodate 10,000 active users. Table 7 describes the detailed building block profile in this solution.

Table 7. Building block profile

Item	Value
Number of Exchange 2013 mailboxes per Exchange Mailbox server (building block)	10,000 (5,000 active and 5,000 passive in a normal scenario)
Mailbox IOPS profile (in DAG configuration)	0.1 IOPS per mailbox
Mailbox size	Start at 1 GB, grow to 2 GB
Target average message size	75 KB
Database read/write ratio	3:2 in a DAG configuration
Deleted items retention (DIR) period	14 days
Log protection buffer (to protect against log truncation failure)	3 days
BDM configuration	Enabled as 24 x 7
Number of databases per Exchange Mailbox server	16
Number of mailboxes per database	625

The next step is to calculate the disk requirements for both the IOPS performance and capacity requirements.

Different disk types generate different numbers of Exchange 2013 IOPS. EMC recommends using the values in Table 8 when calculating the Exchange 2013 IOPS requirements for deployment on VMAX storage arrays.

Table 8. Exchange 2013 IOPS for various disk types on EMC VMAX storage

Disk type	Exchange 2013 random IOPS per disk
Flash	1,250
15 k rpm SAS/FC	180
10 k rpm SAS/FC	130
7.2 k rpm SATA	60

When calculating the disk requirements to support the 10,000-user building block, consider the following criteria:

- Always calculate IOPS performance requirements first, then capacity requirements.
- When calculating IOPS requirement, EMC recommends that you add 20 percent for overhead.
- The capacity requirement includes starting with a user mailbox size of 1 GB with the ability to seamlessly grow to 2 GB. You can accomplish this by using the VMAX 20K Virtual Provisioning feature. Calculate the disk number based on the initial mailbox size, but provision volume size based on the maximum future mailbox size.
- This solution uses 2 TB, 7.2k rpm SATA disks.
- On VMAX systems, we recommend sharing database and log volumes across the same disks, but separating them into different LUNs on the hosts.

To calculate the number of disks required for IOPS performance, use the following formula:

$$\text{Disks required for Exchange IOPS} = \frac{\text{Exchange IOPS requirement} * (\%R + WP * \%W)}{\text{Exchange random IOPS per disk}}$$

This formula includes the following values:

- **Exchange IOPS requirement**—number of Exchange 2013 I/O operations per second, plus any overhead
- **%R**—percentage of I/O for reads
- **%W**—percentage of I/O for writes
- **WP**—RAID write penalty multiplier (RAID 1=2, RAID 5=4, RAID 6=6)
- **Exchange random IOPS per disk**—refer to Table 8 for detailed information.

In the building block design, the Exchange IOPS requirement is 1,000 (10,000 mailboxes * 0.1 mailbox IOPS profile) plus 20 percent overhead. The Exchange 2013 read/write ratio on this mailbox IOPS profile is 60 to 40 percent (3:2). So if we choose RAID 1 for 7.2 k rpm SATA disks:

$$\text{Disks required for Exchange IOPS} = 1000 * (1 + 20\%) * (60\% + 2 * 40\%) / 60 = 28 \text{ disks}$$

After determining the IOPS requirement, the next step is to calculate the capacity requirement. When you use Virtual Provisioning, you must perform two sets of calculations for determining the disk capacity requirements:

- The first set of calculations determines the initial capacity requirements. This identifies the storage requirements needed to support the initial mailbox capacity. In this solution, the initial mailbox capacity is 1 GB.
- The second set of calculations determines thin-provisioned capacity requirements. This is necessary to properly configure the size of the database and log LUNs to be presented to the host. This is also necessary for provisioning the required storage for a fully-provisioned mailbox. In this solution, the full mailbox capacity is 2 GB.

When calculating the capacity requirement, you need to consider not only the mailbox size quota, but also other sizing aspects like database white space, recoverable items folder size, index space, and LUN free space. For more detailed information about building block design methodology and calculation steps, refer to the *Microsoft Exchange Server Best Practices and Design Guidelines for EMC Storage White Paper* on the [EMC website](#).

Microsoft and EMC provide tools to help you properly size your Exchange Mailbox server. The [Exchange 2013 Mailbox Server Role Requirements Calculator tool](#) from Microsoft provides computing and storage recommendations for Exchange sizing. [EMC Exchange 2010-2013 Designer](#) enhances Microsoft's calculator and adds special EMC considerations and extensions.

Table 9 details the IOPS and capacity requirements for the different RAID types.

Table 9. Exchange 2013 IOPS and capacity requirements based on different RAID types

RAID type	IOPS requirement	Capacity requirement	Best option
RAID 1	28 disks	28 disks	28 disks using RAID 1
RAID 5 (3+1)	44 disks	20 disks	
RAID 6 (6+2)	60 disks	18 disks	

This table shows that to support 100,000 total users (ten building blocks), we need 280 * 2 TB, 7.2 k rpm SATA disks in a RAID 1 configuration. Table 10 describes the detailed disk number and LUN size we configured for the tests.

Table 10. Final disk and LUN configuration

Item	Value
Disk type	2 TB 7.2 k rpm SATA disks
Number of disks per building block	28
Number of disks for 10 building blocks	280
RAID configuration	RAID 1
Database LUN size	2600 GB
Number of database LUNs per Exchange Mailbox server	16
Log LUN size	140 GB
Number of log LUNs per Exchange Mailbox server	16

To deploy Exchange 2013 with a heavier mailbox IOPS profile, you can add a few SAS or FC drives and use FAST VP to meet the IOPS performance requirement. In this solution we tested FAST VP on a 0.14 mailbox IOPS profile. [FAST VP design and configuration](#) on page 22 has more details.

Phase 3: Validate the design

You must verify the Exchange 2013 storage design for expected transactional IOPS before placing it in a production environment. To ensure that the environment functions appropriately, EMC recommends you use Microsoft Exchange Server Jetstress 2013 to verify your Exchange storage design. [Solution Validation and Test Results](#) on page 28 provides more information.

Exchange storage design best practices on VMAX

To achieve the best storage performance with Exchange 2013 on EMC VMAX storage, follow these additional general storage design best practices:

- To ensure the best performance and ease of management, isolate the Microsoft Exchange workload (database and logs) from other I/O-intensive applications and workloads by assigning the Exchange workload to its own set of disks. This simplifies troubleshooting of storage-related performance issues. The exception to this guideline is the use of Symmetrix VMAX FAST VP.
- We recommend that you use Virtual Provisioning for Exchange on VMAX systems.

Note: To use Virtual Provisioning for Windows Server 2012, we recommend that you install the Microsoft hotfix [that improves cloud service provider resiliency in Windows Server 2012](#)

- On VMAX storage array, we recommend that you share database and log volumes across the same disks, but separate them into different LUNs on the hosts.
- In a mailbox resiliency configuration, deploy each DAG copy on its own set of physical disks.

- Use striped meta volumes for optimal performance of Exchange 2013.
- Spread the load as evenly as possible across VMAX storage array resources, including front-end processor and ports, back-end disk resources, and so on. For example, use port 0 of the given front-end processors before using port 1 of the same processors.
- When considering FAST VP, evaluate your current Exchange configuration to identify hot spots and ensure that FAST VP can benefit your design.
- When using FAST VP, exclude Exchange transaction log volumes from the FAST VP policy or pin the log volumes into the tier on which they are created.
- We recommend you select “Allocate by FAST Policy” when configuring FAST VP. This ensures that FAST VP uses all tiers for new allocations based on the performance and capacity restrictions.
- In SAN environments, use redundant host bus adapters (HBAs) connected to different fabrics.
- Install EMC PowerPath/VE on vSphere ESXi hosts for optimal path management and maximum I/O performance.

Note: For more information about installing and configuring EMC PowerPath/VE, refer to the *EMC PowerPath/VE for VMware vSphere Installation and Administration Guide* on [EMC Online Support](#).

- Format Windows NTFS (or ReFS on Windows Server 2012) volumes used for Exchange databases and logs with an allocation unit size of 64 KB.

Note: If using ReFS on Windows Server 2012, you must install the Microsoft hotfix for [Exchange Server 2013 databases become fragmented in Windows Server 2012](#) to correct heavy defragmentation of Exchange databases.

Front-end port usage and zoning

As displayed in Figure 2, the Exchange Mailbox server virtual machines are evenly deployed on the VMware ESXi hosts in this solution. To balance the load across VMAX storage array resources, all the ESXi hosts running Exchange virtual machines are segregated to use different VMAX front-end ports, using the SAN zoning and masking view feature in VMAX. Figure 3 shows the front-end port usage on VMAX for the Exchange workloads.




Figure 3. ESXi hosts HBAs and VMAX front-end ports topology

Thin pool design and configuration

Symmetrix Virtual Provisioning

One of the biggest challenges for Exchange administrators is balancing the provisioning of storage space in the data center to support user mailbox quotas. Administrators typically must allocate all mailbox space upfront based on anticipated storage growth. They do this to reduce the management expense and application downtime incurred when they need to add more storage as their business grows. This generally results in the overprovisioning of Exchange storage capacity, which then leads to higher costs and an increase in power, cooling, and floor space requirements, and lower capacity utilization rates. Even with careful planning, it might be necessary to provision additional storage in the future, which could require Exchange downtime.

Virtual provisioning or thin provisioning enables organizations to reduce costs by increasing capacity utilization for Exchange environments. 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 the regular host-accessible Symmetrix devices. Unlike regular Symmetrix devices, however, 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 thin devices comes from a shared “thin pool” that has been associated with the thin device.

Symmetrix Virtual Provisioning presents a certain amount of virtual capacity to a host, allowing the application to consume space as needed. This technology lowers the

total cost of ownership (TCO) by reducing initial allocation of storage capacity and simplifies storage management by reducing administrative tasks that support growth.

For more information about Symmetrix Virtual Provisioning, refer to the *Implementing EMC Symmetrix Virtual Provisioning with VMware vSphere White Paper* on [EMC Online Support](#).

Thin pool design

EMC Virtual Provisioning greatly simplifies the storage design and helps ensure a balanced configuration across all resources. In this solution, we created one thin pool on VMAX 20K consisting of SATA disks to house all the Exchange database and log volumes. Table 11 details the SATA thin-pool configuration.

We also created a thin pool of FC disks to test FAST VP as described in [FAST VP design and configuration](#).

Table 11. SATA thin pool configuration

Thin pool name	Drive type	RAID type	No. of drives	Thin device size	No. of Thin devices	Purpose
SATA pool	2TB SATA 7.2k rpm	RAID 1	280	2,600 GB	160	Database volumes
				140 GB	160	Log volumes

FAST VP design and configuration

Symmetrix FAST VP

In this solution, we tested FAST VP by adding a FC tier to the existing SATA tier to validate the improvement to Exchange 2013 IOPS performance.

Symmetrix FAST VP operates on Virtual Provisioning thin devices and uses intelligent algorithms to continuously analyze devices at the sub-LUN level. FAST VP identifies and relocates the specific parts of a LUN that are most active and benefit when moved to higher-performing storage, such as flash storage. FAST VP also identifies the least active parts of a LUN and relocates that data to higher-capacity, more cost-effective storage such as SATA, without altering performance.

FAST VP bases automatic and nondestructive data movement between tiers on performance measurements and user-defined policies.

For FAST VP to operate, you need to configure the following three storage elements:

- **Storage tiers**—Resource specifications for the disk technology type (Flash, FC, or SATA), which includes the RAID protection type (RAID 1, RAID 5, or RAID 6).
- **Storage groups**—Logical collection of VMAX devices that are managed together.
- **FAST VP policies**—A policy that groups between one and three tiers for which you assign an upper usage limit for each tier. The upper limit specifies the maximum capacity that a storage group associates with the policy while residing on that particular tier.

The *FAST VP for EMC Symmetrix VMAX Theory and Best Practices for Planning and Performance White Paper* on [EMC Online Support](#) provides more information about FAST VP.

FAST VP design

Due to the changes in the Exchange 2013 storage architecture, which has resulted in lower I/O to storage devices and the trend to deploy larger mailboxes, many Exchange designs are capable of utilizing high capacity low rpm disks (for example, 7.2 k rpm SATA).

Whether to use FAST VP with Exchange on VMAX storage arrays depends entirely on customer requirements. If the customer needs FAST VP for all applications or needs to be able to handle unanticipated spikes in performance demand, FAST VP is a good fit. FAST VP reduces administration workloads and provides faster space and I/O issue resolution.

In this solution, an additional FC tier was introduced to meet the requirement of higher demand for Exchange I/O. We considered the following aspects for Exchange application workload for FAST VP:

- When sizing FAST VP tiers, ensure the skew (working set) for the application will fit with room to avoid thrashing between tiers. Exchange's working set is typically around two to five percent.
- For the Exchange workload, exclude transaction log volumes from the FAST VP policy or pin all the log volumes into the tier on which they are created.
- Select **Allocate by FAST Policy** to use FAST VP for all tiers and new allocations, which results in performance and capacity restrictions. This is a new feature introduced in Enginuity 5876 code.
- In this solution, we chose a RAID 5 protection type for FC tiers to yield the best TCO and RAID 1 mirrored protection for SATA to yield the best performance results.

Table 12 lists the FAST VP tiers and disk information for this solution.

Table 12. FAST VP design

Tier	Number of disks	RAID type	FAST VP policy
FC: 600 GB, 10k rpm	40	RAID 5 (3+1)	3%
SATA: 2 TB, 7.2k rpm	280	RAID 1	97%

Using Unisphere to configure FAST VP

Unisphere for VMAX provides features and functionality for managing FAST VP, where all FAST VP managed objects and associations, configuration settings, and time windows can be created and managed.

Complete the following steps to create the Tiers:

1. In Unisphere, click the **Storage** tab and select **Tiers**.
2. Click **Create** and type the applicable tier information, as shown in Figure 4. Then click **OK** to create the tier.




Figure 4. Creating tiers in Unisphere for VMAX

Complete the following steps to enable FAST VP and create the FAST Policies:

1. In Unisphere, click the **Storage** tab and select **FAST**.
2. Click **Settings** and then click **Enable** to set FAST VP, as shown in Figure 5. You can also click either **Performance Time Window** or **Move Time Window** to set up the FAST VP time windows. Click **OK**.




Figure 5. Enabling FAST VP in Unisphere for VMAX

3. Click **Manage policies** and then click **Create** to create a new FAST Policy.
4. In **Create FAST Policy**, type the **Policy Name**, pick the **Emulation**, select the desired **Tiers**, type the maximum percentage of storage group that should reside in each tier, as shown in Figure 6, and then click **OK**.

Create FAST Policy

* Policy Name: R5FC_R1SATA

* Emulation: FBA

Tier *	RAID5_FC_VP	3	%
Tier	RAID1_SATA_VP	97	%
Tier	N/A		%
Tier	N/A		%

OK Cancel Help

Figure 6. Creating FAST Policy in Unisphere for VMAX

5. Select the new FAST VP policy and click the **Associate Storage Group** to select storage group for the policy.

Implementing Fully Automated Storage Tiering for Virtual Pools (FAST VP) for EMC VMAX Series White Paper on [EMC Online Support](#) has more detailed information on how to configure FAST VP in Unisphere.

Using VSI for storage provisioning

You can use EMC Virtual Storage Integrator (VSI) for VMware vSphere to provision VMAX storage for VMware ESXi hosts.

VSI for VMware vSphere is a plug-in to the vSphere client that provides a single management interface that is used for managing EMC storage within the vSphere environment. Unified Storage Management (USM) is a feature of VSI for VMware vSphere designed to simplify storage administration on EMC storage platforms including VMAX. USM is distributed as a Zip file with a single-file installer that is available for download from [EMC Online Support](#).

Before you can use USM for VMAX systems, confirm the following prerequisites:

- Install EMC SMI-S Provider on your SMI-S server.
- Configure six gatekeepers for each Symmetrix array accessed by the provider on your SMI-S server.

When started, the SMI-S Provider automatically discovers all VMAX systems connected to the array provider host. No other action is required, including the use of SYMCFG to discover commands.

- Set up authentication with the following steps:
 - Go to https://<SMI-S_server_IP_address>:5989/ecomconfig (where <SMI-S_server_IP_address> is your real IP address) and log in with **admin** as the username and **#1Password** as the password.
 - Create a new user with the role of Administrator.

Complete the following steps to provision volumes for ESXi hosts:

1. In the vSphere client, right-click an object (the object can be a host, cluster, folder, or data center), and select **EMC > Unified Storage > Provision Storage**.
2. All storage used by an Exchange guest machine for storage of Exchange data must be block-level storage in the VMware virtualization environment. Select **Disk/LUN**, and then click **Next**, as shown in Figure 7.




Figure 7. Storage type selection in VSI USM

3. Select the applicable VMAX array and click **Next**, as shown in Figure 8.




Figure 8. VMAX array selection in VSI USM

4. Select the applicable thin pool on the VMAX array and click **Next**, as shown in Figure 9.

Storage Pool Name	RAID Type	Disk Type	Available Capacity (GB)	Subscription %
FC pool	RAID5	FC	unlimited	0%
SATA pool	RAID1	SATA	unlimited	209%

Figure 9. Thin pool selection in VSI USM

5. Follow the wizard to choose the masking view and VMFS version, and to create a VMFS datastore or an RDM volume with a specified size. Figure 10 shows how to create an RDM volume of 2,600 GB size.

Volume Type

☐ VMFS Datastore
Datastore Name:

☒ RDM Volume

LUN Properties

Capacity: GB
(Max: 65535.99 GB)

Figure 10. Datastore or volume creation in VSI USM

EMC VSI for VMware vSphere: Unified Storage Management—Product Guide on [EMC Online Support](#) provides more information about the system requirements and instructions for using Unified Storage Management (USM).

Solution Validation and Test Results

Overview

In this solution, we used Microsoft Jetstress to test the I/O throughput on the disk subsystem with required performance constraints of Exchange. Jetstress simulates Exchange I/O at the database level by interacting directly with the Extensible Storage Engine (ESE) database technology (also known as Jet) without requiring Exchange to be installed. To simulate the Exchange I/O accurately, Jetstress uses the same ESE.dll file that Exchange uses in production.

You can configure Jetstress to test the maximum I/O throughput available to the disk subsystem within the required performance constraints of Exchange. Jetstress can accept a simulated profile of specific user counts and IOPS per user to verify that the hardware and software components within the I/O stack, from the operating system down to the physical disk drive, are capable of maintaining an acceptable performance level. You can download Jetstress 2013 from [Microsoft Exchange Server Jetstress 2013 Tool](#) on the Microsoft website.

Before running Jetstress, you must know which key metrics to capture and what thresholds to meet for each test metric. Table 13 lists the key metrics for Jetstress verification.

Table 13. Key metrics for Jetstress verification

Performance counters	Target values
Achieved Exchange transactional IOPS (I/O database reads/sec + I/O database writes/sec)	Number of mailboxes * Exchange 2013 user IOPS profile.
I/O database reads average latency (ms)	Less than 20 ms
I/O database writes average latency (ms)	Less than I/O database reads average latency. This counter isn't a good indicator for client latency because database writes are asynchronous.
I/O log writes average latency (ms)	Less than 10 ms

Notes:

- Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, this workload should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.
 - All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly.
 - EMC Corporation does not warrant or represent that a user can or will achieve similar performance expressed in transactions per minute.
-

Test scenarios

To verify the storage design, we used Jetstress 2013 version 15.00.0658.004 to simulate an I/O profile for each mailbox. We validated the Exchange building blocks using a two-hour performance test. The total 100,000 active users were simulated on ten Mailbox servers (10,000 active users per Mailbox server) to validate the performance under the worst (failover) situation.

We tested the following two scenarios to in this solution:

- Baseline: Jetstress test running on all 280 SATA disks on the storage array. The IOPS target is 10,000 mailboxes * 0.1 IOPS per mailbox = 1,000 IOPS per Mailbox server (10,000 IOPS in ten Mailbox servers).
- FAST VP: Jetstress test with two tiers on the storage array, 280 SATA disks and 40 FC disks. The IOPS target is 10,000 mailboxes * 0.14 IOPS per mailbox = 1,400 IOPS per Mailbox server (14,000 IOPS in ten Mailbox servers).

Baseline test result

Jetstress result

Figure 11 displays the Jetstress performance results in the baseline test. In this baseline test, we achieved 11,377 Jetstress IOPS across ten Mailbox servers, 1,377 IOPS more than the target of 10,000 IOPS. Meanwhile, the database and log response times are under Microsoft's recommended thresholds, described in Table 13.




Figure 11. Jetstress results in the baseline test

Using Unisphere to analyze performance

The EMC Unisphere for VMAX 1.6 release includes Performance Viewer which takes the Unisphere data and provides the reporting and diagnostic tools for offline analysis and planning.

In Unisphere, click the **Performance** tab and then go to either **Monitor** view or **Analyze** view to see the VMAX performance, as shown in Figure 12.




Figure 12. The Performance tab in Unisphere for VMAX

In the **Monitor** view, you can use dashboards to see a collection of charts that you define, or that Unisphere for VMAX provides.

In the **Analyze** view, you can view Symmetrix system data for various collection ranges:

- **Real Time** view collects data between 2 and 5 seconds for a limited group of metrics. The data is available for the previous hour.
- **Diagnostic** view collects data every 5 minutes for root cause analysis. The data is available for the previous 7 days.
- **Historical** view collects data in 15 minute intervals for trending and planning. The data is available for the previous year.

EMC Unisphere for VMAX Product Guide on [EMC Online Support](#) provides instructions for using Unisphere for VMAX.

We used the **Diagnostic** view to analyze the VMAX performance during the Jetstress baseline test. Figure 13 shows the overall VMAX array performance status during the time when we tested Jetstress baseline. We can see from the charts about how many host IOPS were recorded on the VMAX and the read/write ratio. Many other counters are available in the **Kpi** and **All** tabs for analyzing performance.




Figure 13. Using Diagnostic view to analyze VMAX performance

Figure 14 shows the VMAX front-end director status during the Jetstress baseline test. There was an even workload spread across all utilized front end directors on the VMAX.




Figure 14. Front-end director host IOPS status

Figure 15 shows the VMAX back-end director status during the Jetstress baseline test. Similar to the front-end directors, the workload was also evenly distributed across all back-end directors. The VMAX system is not taxed during stress testing and is capable of handling far greater workloads.




Figure 15. Back-end director % Busy status

FAST VP test result Jetstress result

Figure 16 displays the Jetstress performance results in the FAST VP test. In this test, we achieved 14,542 Jetstress IOPS across ten Mailbox servers, which exceeded the target of 14,000 IOPS. The database and log response times were under the Microsoft recommended thresholds described in Table 13.




Figure 16. Jetstress results in the FAST VP test

Performance improvement between Jetstress baseline and FAST VP tests

Figure 17 shows the comparison of the Jetstress results between the baseline and FAST VP tests. Compared to the baseline test, while the database and log response times remain at the same level, we achieved around 28 percent more IOPS in the FAST VP test with the additional 40 FC drives, based on the existing baseline 280 SATA drives.




Figure 17. Jetstress result comparison between baseline and FAST VP tests

Using Unisphere to check FAST VP status

Unisphere for VMAX provides built-in features to configure and monitor FAST VP status. In Unisphere, click the **Storage** tab and then select **FAST** to view the overall FAST VP status, including the **Tiers Demand Report**, as shown in Figure 18, to check the capacity usage for each tier. The purple triangle shows the maximum demand placed on the tier by FAST VP storage groups, the blue area shows the capacity currently used, and the green area indicates the remaining capacity for future use.




Figure 18. Tier demand report for the FC tier

You can also use the **Diagnostic** view to analyze the VMAX performance. As shown in Figure 19, FAST VP was moving hot Exchange data from the SATA tier to the FC tier after we enabled FAST VP and ran Jetstress tests to simulate the Exchange workload.




Figure 19. FAST VP movement status for the FC tier

Conclusion

Summary

This solution demonstrates that organizations can rely on the EMC Symmetrix VMAX system as a viable and trusted array for Microsoft Exchange 2013. Symmetrix Virtual Provisioning provides a simple, non-invasive, and economical way to provide storage for Exchange, and Symmetrix FAST VP provides efficient storage tiering and flexible configurations to meet variable application requirements.

The Unisphere for VMAX management and monitoring interface reduces provisioning complexity for managing and monitoring VMAX storage. The tight integration of EMC Virtual Storage Integrator (VSI) with VMware vCenter and vSphere enables VMware administrators to work more efficiently.

Findings

The key findings of this solution include the following:

- **Proven Exchange 2013 design methodology:** Using a building-block approach, Exchange 2013 combined with the high-performing EMC Symmetrix VMAX storage arrays can be deployed using a modular design, resulting in predictable performance for all Mailbox servers.
- **Validated performance:** Test results show that we can meet the performance requirements of a total of 100,000 Exchange 2013 mailboxes with a heavy mailbox I/O profile on a VMAX 20K array. FAST VP provides intelligent data movements and improves overall performance. VMAX 20K is an excellent platform to house Exchange 2013 mailboxes.
- **Simplified performance monitoring with Unisphere:** Unisphere for VMAX provides an easy, intuitive, and highly efficient mechanism for managing EMC storage and for monitoring the performance status of the array.
- **Efficient storage provisioning by VSI:** This solution shows that VMware administrators can easily and efficiently provision VMAX storage with a few clicks using VSI in the vSphere environment.

References

White papers

The following documents, available from the [EMC Online Support](#) or [EMC.com](#) websites, provide additional and relevant information. If you do not have access to a document, contact your EMC representative.

- *Microsoft Exchange Server Best Practices and Design Guidelines for EMC Storage*
- *Implementing EMC Symmetrix Virtual Provisioning with VMware vSphere*
- *Implementing Fully Automated Storage Tiering for Virtual Pools (FAST VP) for EMC VMAX Series Arrays*
- *FAST VP for EMC Symmetrix VMAX Theory and Best Practices for Planning and Performance*
- *EMC Host Connectivity Guide for VMware ESX Server*

Product documentation

The following product documents provide additional, relevant information.

- *EMC Symmetrix VMAX Family with Enginuity Product Guide*
- *EMC Symmetrix VMAX Family Arrays with Enginuity Operating Environment Release Notes—Release Level 5876.229.145*
- *EMC Unisphere for VMAX Product Guide*
- *EMC PowerPath/VE for VMware vSphere Installation and Administration Guide*
- *EMC VSI for VMware vSphere: Storage Viewer—Product Guide*
- *EMC VSI for VMware vSphere: Unified Storage Management— Product Guide*

Other documentation

The following topics are available on the [Microsoft Support](#) website:

- [Update that improves cloud service provider resiliency in Windows Server 2012](#)
- [Exchange Server 2013 databases become fragmented in Windows Server 2012](#)
- [Offloaded Data Transfers fail on a computer that is running Windows 8 or Windows Server 2012](#)

Note: The links provided were working correctly at the time of publication.
