

Business Continuity for Microsoft Exchange 2010 Enabled by EMC Unified Storage, Cisco Unified Computing System, and Microsoft Hyper-V

A Detailed Review



EMC Information Infrastructure Solutions

Abstract

The solution described in this white paper was designed, built, and tested by EMC in partnership with Microsoft® and Cisco®. This white paper highlights the benefits of a fully virtualized Exchange 2010 environment built on the Microsoft Hyper-V platform with EMC® CLARiiON® Unified Storage, leveraging Virtual Provisioning™, to provide larger mailbox capacity during initial deployment, and Cisco UCS B-Series servers to provide a full base for the entire Exchange Server and networking infrastructure. Testing was conducted at the [Enterprise Engineering Center \(EEC\)](#), Microsoft's state-of-the-art enterprise solution validation laboratory on their main campus in Redmond, Washington.

November 2010

Copyright © 2010 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

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

Part number: H7337.2

Table of Contents

Chapter 1: Introduction	9
Overview	9
Exchange 2010 Tested Solutions	9
Partnership	9
Virtualized Exchange 2010 solutions	9
Supported virtualization platforms	10
Executive summary	10
Overview	10
Business case	10
Key results	12
Solution overview	12
Purpose	12
Scope	13
Audience	13
Chapter 2: Technology and key components.....	14
Overview	14
Exchange Server 2010 with native DAG mailbox resiliency configuration	14
Windows 2008 R2 Hyper-V	15
Technology	15
Microsoft Hyper-V	15
EMC Unified Storage with Virtual Provisioning	15
EMC CLARiiON overview	15
Why use EMC Unified Storage with Microsoft Hyper-V?	17
Advanced CLARiiON CX4-480 capabilities	17
Cisco Unified Systems B-series servers	18
Platform	18
Components	18
Hardware used in this solution	20
EMC Unified Storage CX4-480	20
Cisco Unified Compute System	20
LAN and SAN switches	21
Software used in this solution	22
Chapter 3: Solution Design	23
Overview	23

Solution design methodology	23
Overview	23
Key requirements	24
Multi-site design architecture	25
Mailbox server high availability	25
Site resiliency considerations and design architecture	25
Virtualizing Exchange server roles	27
Exchange deployment building block	27
Identifying CPU and memory requirements	28
Identifying the Exchange user profile type	28
Identifying CPU requirements for Mailbox server failure contingency	29
Identifying virtual processor allocation for VMs running Exchange server roles	29
Validating CPU capacity of Mailbox server for site failure contingency	31
Summarizing CPU and memory requirements for the Mailbox server VMs	31
Identifying memory requirements for Mailbox server	32
CPU and memory requirements for Mailbox server summary	32
Hub Transport and Client Access server role requirements	33
Hub/CAS server role requirements	33
CPU memory requirements for Exchange Hub/CAS roles	33
DAG Members and Hub/CAS server allocation	33
Active Directory Domain Controller requirements	34
Active Directory requirements	34
Server selection	34
Database availability design (DAG) design	35
What is a DAG?	35
Planning your DAG deployment	35
Three DAGs design	35
Placement of the file share witness	37
Chapter 4: Storage Design	38
Overview	38
Exchange Mailbox server storage design with thin provisioning	38
Methodology for sizing Exchange storage	39
Exchange storage design using building block methodology	39
What is a building block?	39
Step 1. Identify user requirements	40
Step 2. Identify Exchange VM requirements	41
Step 3. Identify and calculate storage requirements based on IOPS and capacity	41
Step 4. Finalize the Exchange VM building-block	46
Storage design summary	47
Building block summary	47

Capacity requirements summary	47
Total storage requirements summary	47
Determine the thin pool strategy	48
Mailbox server configuration summary	50
Chapter 5: LAN and SAN Architecture	51
Overview	51
UCS server architecture	51
Fabric Interconnects	51
SAN zoning and LAN connectivity	52
WAN for database replication	52
Network load balancing and namespace planning	54
Overview	54
Exchange RPC Client Access and Address Book services	54
Regional namespace model	55
Network traffic	55
Monitoring the Outlook client network ports	55
Chapter 6: Best Practices Planning	57
Overview	57
Exchange 2010 best practices	57
Optimizing SAN best practices	57
Mailbox server optimization for EMC storage	58
Unified Computing System best practices and optimization	59
Chapter 7: Solution Validation	60
Overview	60
Introduction	60
Contents	60
Methodology and tools	61
Overview	61
Jetstress 2010	61
Loadgen	62
EMC Unisphere	62
Exchange storage validation with Jetstress	64
Overview	64
Test configuration	64
Jetstress test results and CX4-480 performance	64
Jetstress performance metrics and test results summary	66
Database seeding performance	67

Overview	67
Database seeding considerations	67
Performing a seed operation	68
Seeding performance	68
Environment validation with Loadgen	68
Overview	68
Test 1 – Normal operating condition – peak load	70
Test 1 objectives	70
Test 1 configuration	70
Test 1 performance results and analysis	71
Test 2 – Within a site switchover and host failure – peak load	72
Test 2 objectives	72
Test 2 configuration	72
Test 2 performance results and analysis	72
Test 3 – Site failure simulation	74
Test 3 objectives	74
Test 3 configuration	74
Test 3 performance results and analysis	74
Datacenter switchover validation	76
Datacenter switchover process	76
Step 1. Terminating a partially failed datacenter (DAG is in DAC mode)	76
Step 2. Validate and confirm the prerequisites for the second datacenter	77
Step 3. Activating Mailbox servers (DAG is in DAC mode)	77
Validating primary datacenter service restoration (failback)	78
Overview	78
Mailbox server role failback	78
Chapter 8: Conclusion	80
Appendixes	81
Appendix A: References	81
White papers	81
Product documentation	81
Other documentation	81
Appendix B: Terminology	82

List of figures

Figure 1.	Dynamic capacity management using EMC Virtual Provisioning functionality	11
Figure 2.	Solution logical DAG design diagram	26
Figure 3.	Exchange deployment building block.....	27
Figure 4.	Database distribution per Exchange Mailbox role VMs	49
Figure 5.	Bandwidth requirements from Microsoft's mailbox role calculator	53
Figure 6.	CAS array configuration	54
Figure 7.	Physical architecture for this solution	56
Figure 8.	Unisphere summary page	63
Figure 9.	Exchange 2010 performance on CX4-480 with thin LUNs.	65

List of tables

Table 1.	CLARiiON CX4 storage systems features	16
Table 2.	Cisco Unified Computing System components	18
Table 3.	EMC Unified Storage CX4-480	20
Table 4.	Cisco Unified Compute System	20
Table 5.	LAN and SAN switches	21
Table 6.	Software used in this solution	22
Table 7.	Key customer requirements	24
Table 8.	Message mailbox requirements	28
Table 9.	CPU requirements.....	29
Table 10.	CAS and Hub vCPU requirements.....	30
Table 11.	Megacycles requirements per VM.....	30
Table 12.	Exchange server block megacycle requirements	30
Table 13.	Mailbox server CPU requirements	31
Table 14.	Server memory requirements.....	32
Table 15.	Summary of CPU and memory requirements for mailbox role VM	32
Table 16.	CPU and memory requirements by Exchange role.....	33
Table 17.	DAG implementation	36
Table 18.	Exchange environment requirements	40
Table 19.	Exchange VM requirements.....	41
Table 20.	Mailbox size on disk summary	43
Table 21.	Database capacity requirements.....	44
Table 22.	Database LUN sizes summary.....	44
Table 23.	Log size requirements.....	45
Table 25.	Building-block summary (based on 600 MB initial capacity).....	47
Table 26.	Capacity requirements summary(based on 600 MB initial mailbox capacity).....	47
Table 27.	Thin provisioning storage savings.....	47
Table 28.	Exchange server configurations for this solution	50
Table 29.	Jetstress test results summary.....	66

Table 30.	Seeding performance information	68
Table 31.	Primary counters and validation criteria	70
Table 32.	Environment validation tests with Loadgen.....	70
Table 33.	Validation of expected load for Test 1	71
Table 34.	Performance results for Loadgen Test 1.....	71
Table 35.	Validation of expected load for Test 2.....	72
Table 36.	Performance results for Loadgen Test 2.....	73
Table 37.	Validation of expected load for Test 3.....	74
Table 38.	Performance results for Loadgen Test 3.....	75

Chapter 1: Introduction

Overview

This chapter introduces the contents of this white paper. It includes the following topics.

Topic	See Page
Exchange 2010 Tested Solutions	9
Executive summary	10
Solution overview	12

Exchange 2010 Tested Solutions

Partnership

This paper provides a solution designed by EMC in partnership with Microsoft and Cisco as part of the Exchange 2010 Tested Solutions venture.

Exchange 2010 Tested Solutions is a joint venture between Microsoft and participating server, storage, and network infrastructure partners to examine common customer scenarios and key design decision points facing customers who are planning to deploy Exchange 2010. Through a series of solution white papers, this initiative provides examples of well-designed, cost effective, Exchange 2010 solutions deployed on the latest and greatest available hardware configurations offered by server and storage partners.

Virtualized Exchange 2010 solutions

As part of this new venture, EMC in partnership with Microsoft and Cisco have designed, built, and validated Exchange 2010 solutions that can help customers make decisions about deploying virtualized Exchange 2010 in their environment. The solution described in this white paper demonstrates how deploying Exchange in a virtualized environment can help customers realize the long-term benefits of their server and storage infrastructure investments.

Supported virtualization platforms

Leveraging a Hyper-V virtualization platform with high-performance UCS servers and shared SAN resources provides greater user resource consolidation with more flexible disaster recovery choices. Today, Microsoft fully supports Exchange on both the Hyper-V platform and VMware technology, as well as on all virtualization products that comply with the Microsoft Server Virtualization Validation Program (SVVP). For Exchange Server 2010, Microsoft supports all server roles in virtualized environments, except for the Unified Messaging role.

For more details about the:

- SVVP program, visit Microsoft at <http://www.windowsservercatalog.com/svvp.aspx>
 - Exchange 2010 requirements in virtualization deployments, visit <http://technet.microsoft.com/en-us/library/aa996719.aspx>
-

Executive summary

Overview

In recent years, many IT organizations have been looking for effective cost savings measures for running their IT infrastructure. One of the technologies that allows enterprises to accomplish this goal is virtualization.

Virtualization allows companies to lower IT costs, increase service levels, and add flexibility and agility to their application environment.

This document presents a solution that validates a fully virtualized Exchange 2010 environment built with Microsoft Hyper-V with Cisco UCS B-Series Blade servers with Cisco Nexus 5000 Unified Fabric switches and MDS 9134 Fibre Channel switches connected to EMC's Unified Storage CX4-480 array. This solution showcases the new Exchange native data protection feature.

The solution team (EMC, Cisco, and Microsoft) used comprehensive validation with tools, including Jetstress and Loadgen, to determine the performance boundaries and limits of the entire infrastructure. The team used and documented general best practices to provide information on how to deploy Exchange 2010 across multiple sites.

These guidelines describe how to leverage thin provisioning with EMC storage to provide users with large mailboxes, allowing those mailboxes to grow in size over time without worrying about storage capacity limitations.

Business case

As customers transition from older Exchange Server versions to Exchange Server 2010, they are looking to take advantage of compelling new Exchange features and to improve resource utilization and operational efficiencies. Server virtualization, combined with innovations in server architecture and efficient storage arrays, are among the best methods to achieve the goals of modern IT groups.

Innovations in server architectures enable unprecedented levels of host application resource consolidation enabled by:

- Increased processor speeds
- Greater network and storage capacity
- A reduction in datacenter space, cabling, power, and cooling resources

Along with the server and application consolidation challenges, many organizations see a business need to improve user work efficiency by allowing them to keep older e-mail in their mailboxes for fast retrieval. This presents a challenge because having more e-mails in the messaging environment means that each user needs a much larger mailbox. Larger mailboxes mean that administrators need to allocate more storage to the users in order to satisfy these new business requirements.

Some studies conclude that 31 to 50 percent of allocated storage for messaging infrastructure is unused. Moreover, for moderate Exchange users, it could take years before all these large mailboxes fill to their capacity. During this time, the storage will become cheaper; allowing companies to purchase it at a much lower price than it would be at the time of deployment.

Technology advances like Virtual Provisioning™ (also known as “thin provisioning”) can provide a solution for this challenge. With thin provisioning, companies can provide large mailboxes to their users at the time of deployment, while purchasing the storage necessary to satisfy only the initial capacity requirements as illustrated in Figure 1.

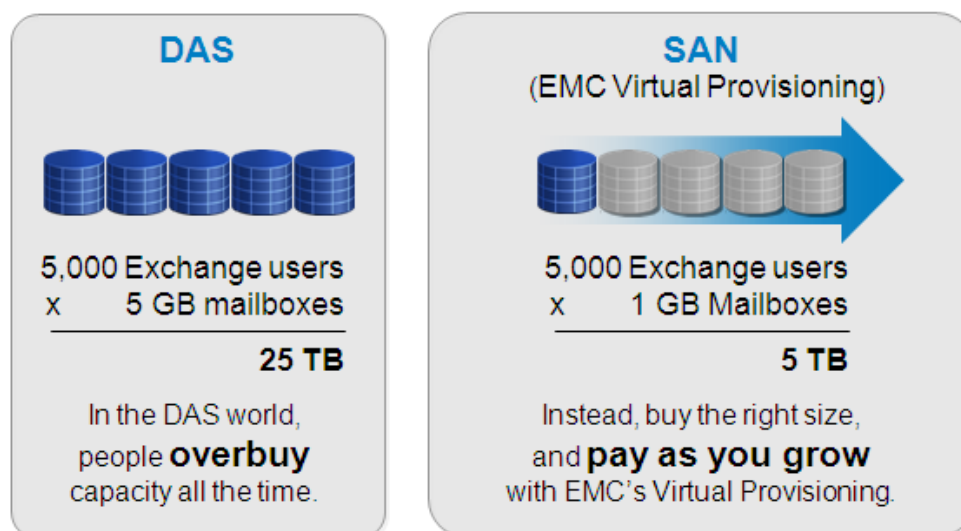


Figure 1. Dynamic capacity management using EMC Virtual Provisioning functionality

Specifically, EMC storage arrays, like EMC Unified Storage and Symmetrix® with Virtual Provisioning technologies, provide additional benefits beyond traditional “thin provisioning” including simplified storage management and improved capacity utilization. Virtual Provisioning enables customers to present a large amount of storage capacity to a host, and then consume only the needed space from a shared pool. When more capacity is required, customers can add it without having to add new databases, modify database copy layouts, or redistribute mailboxes to newly added databases. This lowers total cost of ownership (TCO) by reducing the initial over allocation of storage capacity and simplifies management by reducing the steps required to support growth.

Key results

The solution described in this white paper shows the value of deploying Exchange 2010 on a virtualized environment by consolidating more than 30,000 users onto 12 Cisco UCS B-Series servers across three physical sites. Virtualization in this scenario provides a three to one consolidation ratio that could require up to 36 servers in a physical Exchange deployment.

The tests used in this solution reveal that by using thin provisioning provided by EMC Unified Storage platforms, customers can save up to 243 TB of storage at the time of their Exchange 2010 deployment. EMC thin provisioning allows customers to plan for larger mailboxes for each user during the initial solution deployment. When more storage is required in the future, customers can simply add more disks to the storage pools and grow their database volumes. This procedure does not require downtime and does not affect performance.

Solution overview

Purpose

The purpose of this white paper is to provide customers with information about how to deploy, test, and validate a virtualized Exchange 2010 solution on Microsoft's Hyper-V platform, Cisco's UCS systems, and EMC Unified Storage. The paper includes details about the solution's design and provides performance results recorded during the validation phase.

This white paper describes how to simplify Exchange virtual deployment by leveraging the EMC building-block approach, which helps to deploy virtualized Exchange solutions more easily and effectively. It also describes how to leverage Unified Storage thin provisioning when deploying larger mailboxes in your Exchange Server 2010 environment.

This document's objectives are to:

- Provide guidance and best practice methodologies for designing virtualized multi-site Exchange 2010 solutions.
- Explain how to design an EMC Unified Storage CX4-480 storage array to support thin provisioning for Exchange 2010 mailbox storage.
- Describe how to design and configure Exchange 2010 database availability groups (DAGs) across multiple sites.
- Demonstrate how to design and configure Cisco UCS B-series servers and SAN infrastructures to support a multisite Exchange 2010 deployment.
- Validate the solution using Microsoft tools such as Jetstress and Loadgen.
- List the steps and procedures for performing a switchover and failover.

Scope

The scope of this paper encompasses the following:

- The steps outlined during each stage are high-level in nature. You should read this information in conjunction with the Exchange 2010 documentation referenced at the Microsoft Technet website.
- The setup of this environment simulates an enterprise Exchange 2010 environment. Actual customer configurations will be different. Use the information contained in this white paper only as a part of a rigorous analysis and design exercise.
- Before implementing a solution in a production environment, consider the size and complexity of the environment. EMC recommends that you consult with the EMC Consulting or Microsoft Solutions consultants (MSCs) for onsite assistance with planning, installation, and integration requirements.

The information contained in this document is not intended to replace existing, detailed product implementation guides for deploying UCS servers, SANs, Exchange, or storage infrastructures.

Audience

This white paper is intended for information technology professionals who are involved in the evaluation, architecture, deployment, and data management of data center computer systems, storage systems, and Microsoft Exchange infrastructures.

Chapter 2: Technology and key components

Overview

This chapter identifies and briefly describes the technologies and components used in this solution. It contains the following topics.

Topic	See Page
Exchange Server 2010 with native DAG mailbox resiliency configuration	14
Windows 2008 R2 Hyper-V	15
EMC Unified Storage with Virtual Provisioning	15
Cisco Unified Systems B-series servers	18
Hardware used in this solution	20
Software used in this solution	22

Exchange Server 2010 with native DAG mailbox resiliency configuration

Microsoft Exchange Server 2010 is an enterprise e-mail and communication system that allows businesses and customers to collaborate and share information. EMC enhances Exchange Server 2010 with the industry's broadest choice of storage platforms, software, and services.

With the new version of Exchange 2010, Microsoft presents a new, unified approach to high availability (HA) and disaster recovery (DR) by introducing features such as DAGs, and online mailbox moves. Customers can now implement mailbox servers in mailbox resiliency configurations with database-level replication and switchover.

Major improvements with the application database structure and I/O reduction include support for a larger variety of disk and RAID configurations including:

- High-performance Enterprise Flash Drives (EFDs)
- Fibre Channel (FC) drives
- Slower performing serial advanced technology attachment (SATA) drives
- Serial attached SCSI (SAS) drives

Windows 2008 R2 Hyper-V

Technology Hyper-V is Microsoft's hypervisor-based virtualization technology that is integrated into most Windows Server 2008 x64 operating systems (Hyper-V is not available in the Web or Foundation editions of Windows Server 2008). As a virtualization solution, Hyper-V enables users to take maximum advantage of the server hardware by providing the capability to run multiple operating systems (on virtual machines) on a single physical server.

Microsoft Hyper-V Microsoft Hyper-V is virtualization software that allows you to consolidate your servers by running (as virtual machines) several instances of similar and dissimilar operating systems on one physical machine. This cost-effective, highly scalable, virtual machine (VM) platform offers advanced resource management capabilities. Hyper-V minimizes TCO for your environment by:

- Increasing resource utilization
- Decreasing the number of servers and all associated costs
- Maximizing server manageability

For more details see the following websites:

- Microsoft Hyper-V visit Microsoft technical library at <http://www.microsoft.com/windowsserver2008/en/us/hyperv.aspx>.
 - Microsoft's Exchange 2010 Systems requirements for hardware virtualization visit Microsoft at <http://technet.microsoft.com/en-us/library/aa996719.aspx>.
-

EMC Unified Storage with Virtual Provisioning

EMC CLARiiON overview The EMC CLARiiON[®] family of networked storage systems brings high performance to the mid-tier market with a wide range of storage solutions—all based on the powerful, proven, eight generations of CLARiiON architecture. CLARiiON provides multiple tiers of storage (enterprise Flash drives (EFDs), FC (FC), and SATA) in a single storage system. This system significantly reduces acquisition costs and management costs by allowing users to manage multiple storage tiers with a single management interface.

The next-generation CLARiiON systems, called the CX4 series with UltraFlex[™] technology, deliver storage systems that you can easily customize by populating your I/O slots with either FC or iSCSI I/O modules. Products with multiple back ends such as the CX4-240, CX4-480, and CX4-960 can support disks operating at both two Gb/s and four Gb/s simultaneously.

CLARiiON storage systems address a wide range of storage requirements by providing flexible levels of capacity, functionality, and performance. The CX4-120 supports up to 120 drives and connectivity for up to 128 HA hosts. The CX4-240 storage system expands the family, supporting up to 256 HA hosts and up to 240

drives. The CX4-480 further expands the CX4 family by supporting 256 HA hosts and 480 drives. The high-end CX4-960 adds even more capability, supporting up to 512 HA hosts and up to 960 drives. Table 1 summarizes the basic features for the CLARiiON CX4 storage system.

Table 1. CLARiiON CX4 storage systems features

Feature	CX4-120	CX4-240	CX4-480	CX4-960
Maximum disks	120	240	480	960
Storage processors (SP)	2	2	2	2
Physical memory per SP	3 GB	4 GB	8 GB	16 GB
Max write cache	600 MB	1.264 GB	4.5 GB	10.764 GB
Max initiators per system	256	512	512	1024
High-availability hosts	128	256	256	512
Minimum form factor size	6U	6U	6U	9U
Maximum standard LUNs	1024	1024	4096	4096
SnapView™ snapshots	Yes	Yes	Yes	Yes
SnapView clones	Yes	Yes	Yes	Yes
SAN Copy™	Yes	Yes	Yes	Yes
MirrorView™/S	Yes	Yes	Yes	Yes
MirrorView/A	Yes	Yes	Yes	Yes
RecoverPoint/S	Yes	Yes	Yes	Yes
RecoverPoint/A	Yes	Yes	Yes	Yes

Why use EMC Unified Storage with Microsoft Hyper-V?

CLARiiON and Hyper-V work well together. Some of the reasons CLARiiON is an ideal fit for Hyper-V in the midrange storage market include:

- CLARiiON storage systems provide several flexible levels of models with FC and iSCSI interfaces. This allows the user to make the optimal choice of a storage system based on capacity, performance, and cost.
- CLARiiON storage systems can scale quickly to manage anticipated data growth, especially as the storage needed for VMs increases on Microsoft Hyper-V Server.
- CLARiiON Virtual Provisioning (thin provisioning) improves storage capacity utilization and simplifies storage management by presenting a VM with sufficient capacity for an extended period of time. With CLARiiON Virtual Provisioning, customers can provision less storage. Rather than buying additional capacity up front, customers can reduce or defer initial capacity requirements. Furthermore, customers can save on acquisition and energy costs by running their systems at higher utilization rates and adding capacity as needed without disrupting applications.
- CLARiiON storage can be shared across multiple Microsoft Hyper-V Servers, allowing storage consolidation to provide efficient use of storage resources. This storage consolidation is valuable for clustering and quick migration.
- VM applications and data on CLARiiON storage systems enhance performance and therefore maximize functionality, reliability, and efficiency of Microsoft Hyper-V Server as opposed to internal server storage.
- The Unisphere™ Manager suite provides web-based centralized control of global disk space, availability, security, quality of service, and replication for VMs provisioned by the CLARiiON storage system.
- The redundant architecture of the CLARiiON storage system provides no single point of failure, thereby reducing application downtime and minimizing business impact for storage upgrades.
- The CLARiiON storage system's modular architecture allows a mixture of EFDs, FC, and SATA drives.

Advanced CLARiiON CX4-480 capabilities

Advanced capabilities start with the scalability to meet both the needs of today and the requirements of tomorrow. EMC CLARiiON CX4-480 models are a low-cost approach to deploying external storage. They provide an economical storage platform for applications such as replication, backup-to-disk, and a variety of data archiving tasks. This model's features include:

- 15 drives per enclosure
- Scaling up to 480 drives through additional expansion enclosures
- Four FC front-end and four FC back-end buses per system
- Choice of additional I/O connectivity with UltraFlex I/O modules (FC or iSCSI)
- Storage for up to 256 hosts

With support for both iSCSI and FC connectivity, the EMC CLARiiON CX4-480 enables organizations to use the appropriate network interconnection for their

environments on the same system. In addition, users also have the option to expand their storage in the future. The EMC CLARiiON CX4-480 arrays support cost-effective, shared storage by using iSCSI interconnection with widely available IP networking components for direct-attach to a network, using conventional Ethernet switches.

EMC CLARiiON CX4-480 arrays, using four Gb/s and one Gb/s iSCSI connections, use modularity with I/O connectivity through UltraFlex to provide cost-effective, direct-attach configurations with a wide range of SAN switch options to create SANs for up to 64 high-availability servers. Each controller supports up to 12 x 4 Gb/s FC front-end ports or 16 x 1 Gb/s iSCSI.

The EMC CLARiiON CX4-480 series delivers functionality that releases the benefits of tiered storage. It is the answer to storage consolidation for heterogeneous environments supporting Windows, Linux, AIX, HP-UX, Solaris, and VMware. The EMC CLARiiON CX4-480 supports FC, SATA II, low-power SATA II, and Flash drives.

Cisco Unified Systems B-series servers

Platform

The Cisco Unified Computing System is a next-generation data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce TCO and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain.

Components

Table 2 describes the main system components of the Cisco Unified Computing System (UCS). Working as a single, cohesive system, these components unify technology in the data center. They represent a radical simplification in comparison to traditional systems, helping to simplify data center operations while reducing power and cooling requirements. The system amplifies IT agility for improved business outcomes.

Table 2. Cisco Unified Computing System components

Component	Description
Compute	The system is based on an entirely new class of computing system that incorporates blade servers based on Intel Xeon 5500, 5600, and 7500 series processors. The blade servers offer patented Cisco Extended Memory Technology to support applications with large datasets and allow more VMs per server.
Network	The system is integrated onto a low-latency, lossless, 10 Gbps unified network fabric. This network foundation consolidates what today are three separate networks: LANs, SANs, and high-performance computing networks. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.

Virtualization	The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
Storage access	The system provides consolidated access to both SAN storage and network attached storage (NAS) over the unified fabric. Unifying storage access means that the Cisco Unified Computing System can access storage over Ethernet, FC, FC over Ethernet (FCoE), and iSCSI, providing customers with choice and investment protection. In addition, administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity and management while helping to increase productivity.
Management	The system uniquely integrates all the system components, so that users are able to manage the entire solution as a single entity through Cisco UCS Manager software. Cisco UCS Manager provides an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application-programming interface (API) to manage all system configuration and operations. Cisco UCS Manager helps increase IT staff productivity, enabling storage, network, and server administrators to collaborate on defining service profiles for applications. Service profiles are logical representations of desired physical configurations and infrastructure policies. They help automate provisioning and increase business agility, allowing data center managers to provision resources in minutes instead of days.

Hardware used in this solution

EMC Unified Storage CX4-480

Table 3 provides information about the EMC hardware used in this solution.

Table 3. EMC Unified Storage CX4-480

Item	Description
Storage	3 EMC Unified Storage CX4-480s – 1 per site
Storage connectivity to host(FC or iSCSI)	FC
Storage cache	16 GB
Number of storage controllers	2 per storage frame
Number of storage ports available/used	8 (4 per system processor (SP)) available per storage frame, 4 used (2 per SP)
Maximum bandwidth of storage connectivity to host	8 * 4 Gb/s
Disk type used in this solution	450 GB 15k rpm, FC
Total number of disks tested in solution	432 (360 for DBs and 72 for logs across 3 sites, 144 per site)
Maximum number of spindles that can be hosted in the storage	480 in a single storage array

Cisco Unified Compute System

Table 4 provides information about the Cisco hardware used in this solution.

Table 4. Cisco Unified Compute System

Item	Description
Blade Server	4 x B200 M1 per site
Processors	2 x Intel Xeon x5570 (2.93GHz) per site
Memory	96 GB RAM (12 x 8GB DIMM) per server
Converged Network Adapter	M71KR-Q (2 x 10 Gigabit Ethernet and 2 x 4 Gbps FC) per server
Internal Blade Storage	2 x 146 GB SAS 10k rpm disk (RAID1) per server
Chassis	5108 (6RU) one per site
Fabric Extender	2 x 2104XP per site
Fabric Interconnect	2 x 6120XP per site
Fabric Interconnect Expansion Module	2 x 8-port 4 Gb/s FC per site

LAN and SAN switches

Table 5 provides information about the LAN and SAN switches used in this solution.

Table 5. LAN and SAN switches

Item	Description
10 Gigabit Ethernet switch	2 x Nexus 5010 (8 fixed 1GbE/10GbE ports, 12 fixed 10GbE ports, Datacenter Bridging) per site
FC switch	2 x MDS 9134 (32 fixed 4 Gbps ports) per site

Software used in this solution

Table 6 provides information about software used in this solution.

Table 6. Software used in this solution

Item	Description
Hypervisor host servers	Windows 2008 R2 Hyper-V Enterprise
Exchange Server VMs	Windows 2008 R2 Enterprise
Exchange Server 2010 Mailbox server role	Enterprise Edition RU3 or later
Exchange Server 2010 Hub Transport and Client Access Server Role	Standard Edition RU3 or later
Multipath and I/O Balancing	EMC PowerPath®
Antivirus Software (on Hub Transport servers)	ForeFront Protection 2010 for Exchange Server (v11)

Chapter 3: Solution Design

Overview

This chapter describes the solution's design and contains the following topics:

Topic	See Page
Solution design methodology	23
Multi-site design architecture	25
Identifying CPU and memory requirements	28
Hub Transport and Client Access server role requirements	33
Active Directory Domain Controller requirements	34
Database availability design (DAG) design	35

Solution design methodology

Overview

The following section describes the methodology used to design this solution. The solution team used customer requirements and design assumptions as the key decision points during the Exchange 2010 environment implementation phase. The following questions represent some of the information the solution team considered during the design phase:

- What is the high-availability strategy for Exchange deployment?
- Is there a business requirement for datacenter and application virtualization?
- Is the current network and storage infrastructure sufficient to support the new Exchange 2010 features for high availability?
- Are the network bandwidth between sites and network latencies within the recommended guidelines to support database seeding and log shipping over the long distance?

Key requirements

Table 7 summarizes the key customer requirements and assumptions upon which the solution team designed, built, and validated this Exchange solution.

Table 7. Key customer requirements

Requirement	Description
32,400 users across three active sites	<ul style="list-style-type: none">• 10,800 active Exchange users per site• 100% user concurrency during normal operation
Exchange User Profile Requirements	<ul style="list-style-type: none">• 2 GB mailbox quota (initial planned capacity of 600 MB per mailbox)• 100 messages sent or received per day for each user profile• 100% Outlook MAPI clients (cached mode)
Virtualization and consolidation for any new server and application in the datacenter	<ul style="list-style-type: none">• Consolidation of all server roles by leveraging virtualization• Virtualization of all Exchange roles
HA requirements	<ul style="list-style-type: none">• The HA must be able to tolerate an individual server failure or undergo maintenance without inducing service degradation or requiring site failover• The HA must be able to tolerate a single site failure without incurring service degradation• In site local recovery time objectives (RTO) and recovery point objective (RPO) requirements—5 minute RTO and 1 MB RPO (1 log file)• Site failure RTO and RPO—1 hour RTO and 5 MB RPO (5 log files)
Storage requirements	<ul style="list-style-type: none">• Existing SAN infrastructure must be leveraged <p>Customer already leveraging SAN with an EMC Unified Storage CX4-480 with 450 GB 15k rpm drives</p> <ul style="list-style-type: none">• Storage must support additional I/O overhead of 50% to support future user load. <p>This additional IOPS overhead is required to support future Exchange acquisitions and to share CX4-480 with other applications.</p>

Multi-site design architecture

**Mailbox server
high availability**

In order to address the server high-availability requirement, two copies of every mailbox database must be available. One copy is active and the other is passive. Both the active and passive copies of each mailbox database must be hosted on separate mailbox servers. If the Mailbox server that hosts the active database fails, or if an active mailbox database becomes inaccessible, the passive mailbox database becomes active on a different server in the same site.

**Site resiliency
considerations
and design
architecture**

The DAG feature provides high availability and site resiliency for Exchange 2010 Mailbox servers. It manages database replication and activation. You can add and manage up to 16 Mailbox servers in the same DAG. DAGs provide server high availability within a site and can span geographical sites to provide site resiliency.

Site resiliency requires an additional passive copy of every mailbox database hosted on a remote site's Mailbox server. When a site failure occurs, the passive mailbox databases at the remote site become active, providing client access to the mailbox content.

When planning a site resiliency strategy for Exchange mailbox databases, you need to consider several conditions to create an effective DAG design, including network bandwidth/latency between sites and desired failover behavior. This solution has active mailbox users in all three locations and therefore the solution should include three separate DAGs with each DAG spanning two sites.

Figure 2 illustrates how this solution distributes the DAGs across the three sites. In a solution where active mailbox users exist in all three sites, it is a best practice not to deploy a single DAG. With active mailboxes in more than one site, we cannot use the database activation coordination (DAC) mode to prevent unwanted database activations in an alternate site. This means that something as simple as a small, short-lived, network connectivity loss could result in an unnecessary database switchover or a quorum loss between DAG members. Deploying three active/passive DAGs will prevent this scenario.

Each DAG has its primary member servers in one site. These primary member servers host the active and passive mailbox database copies for that site. The same DAG also has member servers in one of the other two sites. These additional members host a third (passive) mailbox database copy. The DAGs operate in an active/passive mode from the perspective of site resiliency. The DAG member servers in one site have all of the active mailbox database copies while member servers of the same DAG have the passive mailbox database copies.

DAG design is discussed in more detail later in the paper.

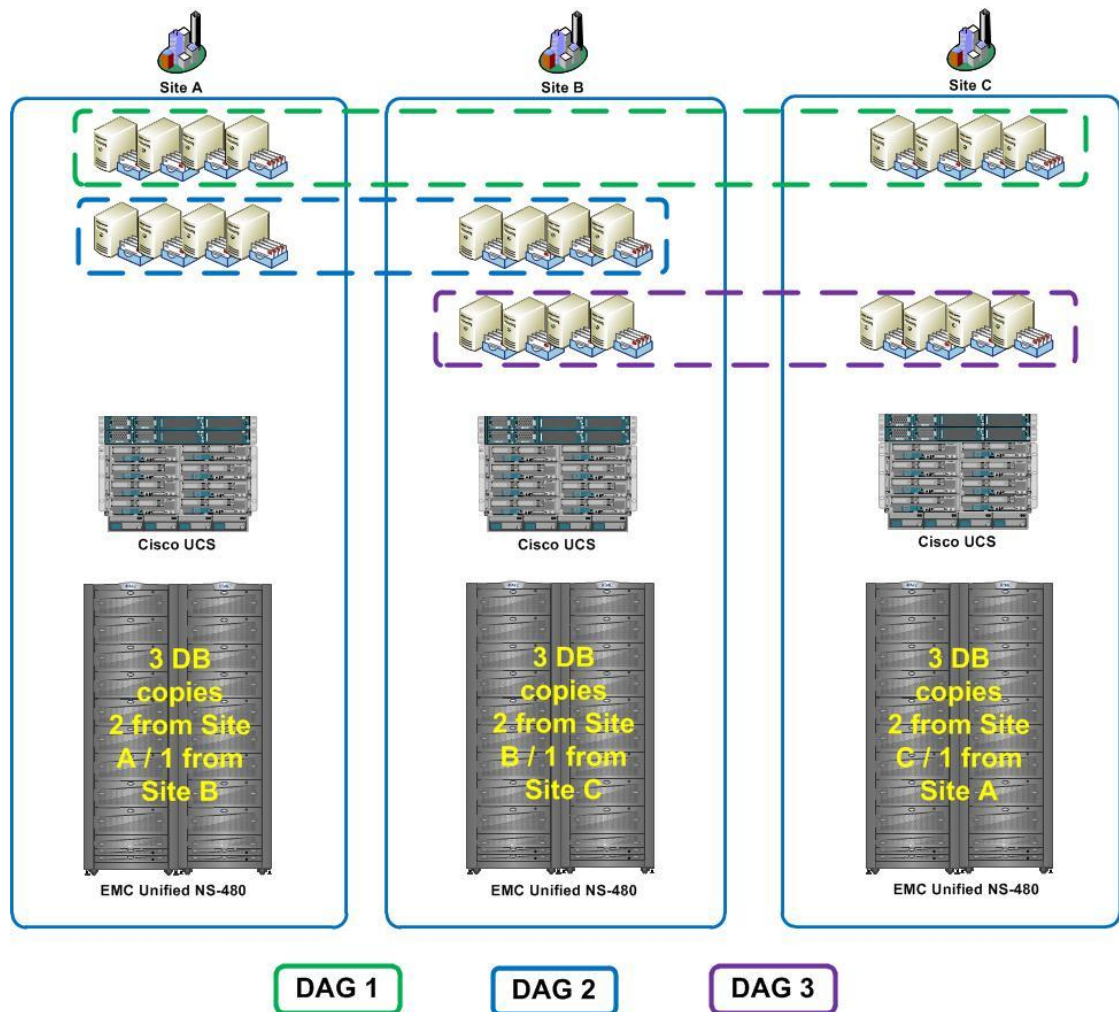


Figure 2. Solution logical DAG design diagram

Virtualizing Exchange server roles

A Mailbox server role running on a physical server can be a member of only one DAG. However, running the Mailbox server in a VM enables a physical server to host mailboxes that belong to different DAGs. The Mailbox server in each VM can be a member of a different DAG.

An individual physical server is a single point of failure. Though the servers have many redundant components, if the physical server experiences a catastrophic failure, all VMs that run on that server will also fail. By placing Mailbox servers with active and passive mailbox databases replicated from another server local site and a second Mailbox server with passive mailboxes replicated from a remote site, none of your users will lose Exchange access if an individual server failure occurs.

The goal is to create a configuration that does not result in users losing access to their Exchange mailbox data. This can be achieved by designing Mailbox servers so that they host both active and passive database copies.

Exchange deployment building block

Exchange 2010 requires three Exchange Server roles in order to deliver functional messaging and collaboration services:

- Mailbox server role
- Client Access server role
- Hub Transport server role

The focus thus far has been on the Mailbox server role because it is the core Exchange role. We have discussed hosting two VMs running the Mailbox server role on the same hypervisor host. Each of the two Mailbox servers is a member of a different DAG.

By adding a third VM that runs the Hub and CAS roles, we have the basis of an Exchange building block that can scale up the deployment by implementing multiple instances to support the required number of mailboxes.

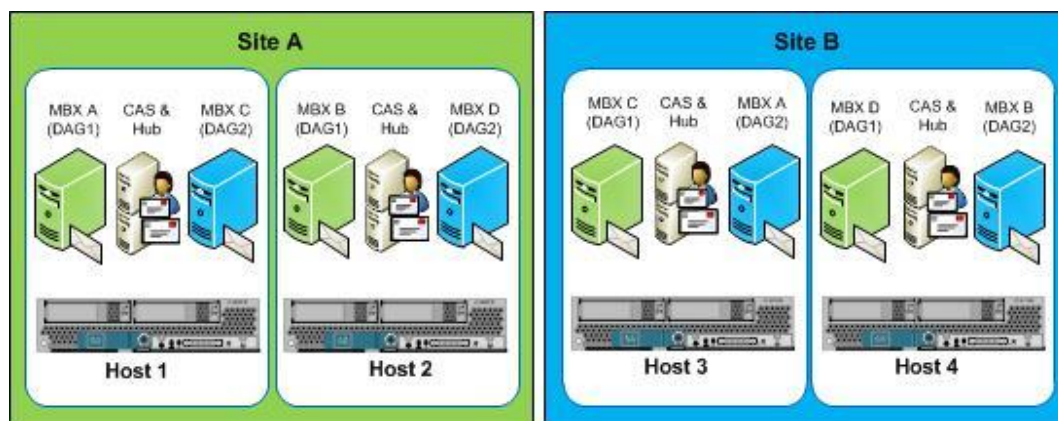


Figure 3. Exchange deployment building block

Identifying CPU and memory requirements

Identifying the Exchange user profile type

To start identifying CPU and memory requirements for this solution we need to know the Exchange user profile type. Customer requirements direct us to size for a 100-message profile. The 100-message mailbox user action profile has the following characteristics and requirements:

Table 8. Message mailbox requirements

Parameter	Value
Messages sent per 8-hour day	20
Messages received per 8-hour day	80
Required megacycles per active mailbox	2
Required megacycles per passive mailbox	0.3
Mailbox replication megacycles factor per database copy	0.1
Required database cache memory per mailbox (MB)	6

The DAG design specifies two Exchange Mailbox server role VMs and one HUB/CAS VM per physical Hyper-V host with eight logical processors. In this design, one Exchange Mailbox server role VM hosts active and passive copies at the local site and the other VM hosts a third passive copy from a remote site.

CPU and memory requirements must support server and site failure contingencies. The servers must be sized to handle the following conditions:

- **A single server failure.** If a single server fails, you need to make sure that there is no service degradation or failover requirements to another site. The failed server can be a Mailbox server, Client Access server, or a Hub Transport server.
- **A site failure.** Accommodating a site failure requires that the surviving site support the failed site's entire workload.

Identifying CPU requirements for Mailbox server failure contingency

CPU and memory requirement calculations start with the Mailbox server role. During a server failure, the surviving servers at a site must support 10,800 mailboxes with a maximum of 100 percent mailbox access concurrency. Additionally, we need to provision sufficient megacycles so that CPU utilization does not exceed 80 percent. The following table lists the Mailbox server megacycle and memory requirement calculations.

Table 9. CPU requirements

Parameter	Value
Active megacycles	10,800 mailboxes x 2 megacycles per mailbox = 21600
Passive megacycles	0 (No passive mailboxes in this situation)
Replication megacycles	0.2 X 21600 active megacycles = 4320
Maximum mailbox access concurrency	1.00 (100%)
Total required megacycles during Mailbox server failure	(21600 + 4320) X 1.00 = 25920
Minimum required server megacycles	25920 / 0.80 = 32400

Cisco UCS B200 M1 servers with two Intel Xeon x5570 CPUs, and 96 GB memory deliver 5,500 new platform mailbox megacycles per CPU core and has a total of eight cores. Cisco obtains this value with Hyper Threading disabled in the testing environment. The total new platform mailbox megacycle capacity for this server is 51,200 (6400 x 8). Based on the version of Windows 2008 R2 with Hyper-V used in this deployment, we can allocate a maximum of four virtual CPUs per guest VM. Hypervisor overhead is estimated at 10 percent for typical deployments. The following formula identifies the number of net megacycles per CPU core available to the VM.

Net New Platform Mailbox Megacycles per Logical Processor Core (LP) = 6400 X (1 – 0.1) = 5500

Identifying virtual processor allocation for VMs running Exchange server roles

The design used for this solution includes three VMs per Hyper-V host. Two of the VMs will run the Mailbox server role and the third VM will run the Hub and CAS role. We also need to determine the number of Exchange Server blocks needed to support the mailboxes in each site to satisfy the site resiliency requirements.

The Mailbox servers that host active mailboxes during normal runtime are paired together to address the server failure and server maintenance requirements. There is one pair of these Mailbox servers in our Exchange Server block design. The second pair of Mailbox servers in our Exchange Server block hosts mailboxes that are replicated from a remote site and are activated when a remote site fails. These Mailbox servers (hosting mailboxes replicated from a remote site) need 50 percent of the capacity that is required by the Mailbox server that hosts local mailboxes because they are not designed to tolerate a concurrent local server failure and remote site failure.

Based on these parameters, we allocate the maximum number of virtual processors to the VMs that host the local mailbox copies. Each of these VMs is allocated four

virtual processors. The VMs that host copies of remote mailboxes are allocated two virtual processors.

The CAS and Hub server combination is typically allocated the same number of virtual processors that are required by the active Mailbox servers in the site. However, we expect the CAS/Hub role to require fewer compute resources than the Mailbox server role. For this reason, we allocate three virtual processors to the CAS/Hub Server role.

Table 10. CAS and Hub vCPU requirements

Virtual machine role	Virtual processors per VMs
Mailbox server (local mailboxes)	4
Mailbox server (replicated mailboxes from the remote site)	2
CAS and Hub combination	3

The next step is to determine if one Exchange Server block is sufficient to support the active mailboxes in a site during various failure and maintenance conditions. The first calculation identifies the available megacycles for each virtual processor, as shown in Table 11.

Table 11. Megacycles requirements per VM

Parameter	Value
Virtual processor megacycles	Net new platform mailbox megacycles per core times the total number of logical processors divided by the total number of virtual processors: $4950 \times (8 / 9) = 4400$

To determine the capacity of the Exchange Server Block, multiply the number of virtual processors in each VM running the mailbox role by the virtual processor's megacycle value, as shown in Table 12.

Table 12. Exchange server block megacycle requirements

Parameter	Value
Megacycles per mailbox VM with 4 virtual processors	$4400 \times 4 = 17600$
Megacycles per mailbox VM with 2 virtual processors	$4400 \times 2 = 8800$

The number of deployed Exchange Server blocks must support all of the mailboxes that are local to a site as well as the mailboxes that are replicated from a remote site and activated during a site failure. The following formula determines the number Exchange Server blocks required to support the minimum number of Mailbox server megacycles for the local site.

$$\text{Number of Exchange Server Blocks} = 32400 / 20480 = 1.84$$

Since more than one Exchange Server block is required, and we cannot have a fraction of an Exchange block, we need two Exchange Server blocks to support the mailboxes. This means we need two VMs to support the mailbox roles.

Validating CPU capacity of Mailbox server for site failure contingency

We now need to validate that the VM in our Exchange Server block meets the site resiliency requirements for our deployment. To meet this requirement, each site needs the capacity to support an additional 10,800 active mailboxes in case its partner site fails. We have already computed the megacycle requirements for 10,800 mailboxes to be 32,400. Additionally, we have already estimated that each of our two logical processor VMs delivers 10,240 megacycles. The following calculation validates that the number of Exchange Server blocks with two logical processor VMs are required for supporting the additional 10,800 mailboxes that will become active during a site failure.

$$\text{Number of Exchange Server blocks} = 32400 / (2 \times 8800) = 1.84$$

Since the result of this equation is less than two, we can conclude that two Exchange Server blocks are sufficient for hosting the mailbox serves with the database copies that are replicated from a remote site and activated during a site failure.

Summarizing CPU and memory requirements for the Mailbox server VMs

We have identified the virtual processor requirements for each VM running the Exchange server roles. Based on the number of mailboxes at each site, we see that the VMs with four virtual processors have a megacycle capacity for supporting 5,400 local mailboxes, while the VMs with two virtual processors support 2,700 mailboxes that are replicated from a remote site and activated during a site failure. Table 13 summarizes this data.

Table 13. Mailbox server CPU requirements

VM Role	CPUs per VM
Mailbox (5,400 Mbx)	4
Mailbox (2,700 Mbx)	2

Identifying memory requirements for Mailbox server

Table 14 shows the mailbox database cache and operating system memory requirement calculations.

Table 14. Server memory requirements

VM hosting 5400 mailboxes	
Parameter	Value
Minimum required mailbox database cache	6 MB x 5400 mailboxes x 1.00 = 31.6 GB
Minimum operating system memory requirement	(Minimum required mailbox database cache x 0.1) + 4 GB = 7.2 GB
Total minimum required memory during Mailbox server failure	31.6 GB + 7.2 GB = 38.8 GB
VM hosting 2700 mailboxes	
Parameter	Value
Minimum required mailbox database cache	6 MB x 2700 mailboxes x 1.00 = 15.8 GB
Minimum operating system memory requirement	(Minimum required mailbox database cache x 0.1) + 4 GB = 5.6 GB
Total minimum required	21.4 GB

CPU and memory requirements for Mailbox server summary

Based on these calculations, we make the following vCPU and memory allocations to the VMs used to run the mailbox roles:

Table 15. Summary of CPU and memory requirements for mailbox role VM

VM Role	vCPUs per VM	Memory per VM
Mailbox (5400 Mbx)	4	39 GB
Mailbox (2700 Mbx)	2	22 GB

Note:

For more details about Exchange 2010 memory and CPU requirements, visit “Understanding Memory Configurations and Exchange Performance” at <http://technet.microsoft.com/en-us/library/dd346700.aspx>.

Hub Transport and Client Access server role requirements

Hub/CAS server role requirements

Now that we have identified the Mailbox role CPU and memory requirements, we can calculate the Hub Transport and Client Access server role (HUB/CAS) requirements. The combined HUB and CAS role has a 1:1 CPU core ratio to the Mailbox server. The Hub/CAS servers must be designed to address the heaviest planned load condition, which is 20,800 active mailboxes during a site failover. Since four vCPUs are allocated to support 5,400 active mailboxes, four vCPUs are allocated to each VM running the Hub/CAS combined role.

Note:

This solution uses Microsoft ForeFront Protection 2010 for Exchange Server software to protect its data. For more information on this product, visit <http://www.microsoft.com/forefront/protection-for-exchange/en/us/>.

CPU memory requirements for Exchange Hub/CAS roles

Memory for the Hub/CAS role is also allocated based on the allocated Mailbox server role CPU core number. Two GB of RAM are allocated for every CPU core. Eight GB of RAM are allocated to the three-vCPU VM that is hosting the Hub/CAS role.

Table 16. CPU and memory requirements by Exchange role

Virtual machine role	vCPU per VM	Memory
Hub/CAS	3	8 GB

DAG Members and Hub/CAS server allocation

These calculations show that each blade server can host two VMs running the Mailbox server role and a third VM that runs the Hub/CAS server role. Placing two B200 blade servers into a local site and a remote site provides a building block that supports 2,700 mailboxes at two sites. Placing eight B200 blades servers into each site provides support for 10,800 mailboxes in each of the three sites for 32,400 mailboxes.

Active Directory Domain Controller requirements

Active Directory requirements

Active Directory server CPU requirements for supporting the Exchange deployment are also calculated as a ratio of Mailbox server role cores to Active Directory cores. The ratio of equivalent CPU cores is 8:1 (Mailbox server role CPU cores to Active Directory server CPU cores).

Note:

An 8:1 ratio assumes 64-bit DC/GC server. For a 32-bit DC/GC server the ratio is 4:1. For additional references visit <http://technet.microsoft.com/en-us/library/dd346701.aspx>.

It is a best practice to accommodate the entire Active Directory database into memory at one time. Four GB to eight GB RAM on a 64-bit Active Directory server enables the caching of extremely large active directory deployments.

For this deployment, the Active Directory services are provided outside the Exchange deployment. The Active Directory team worked closely with the Exchange design team to ensure that the Exchange Active Directory requirements are met for this Exchange deployment. Redundant Active Directory servers, as well as DNS servers, are provisioned in all three sites.

Server selection

A B200 Blade Server with dual Xeon x5570 processors has eight CPU cores, and can support up to 96 GB RAM, using 8 GB DIMMs. This server blade configuration meets the identified requirements for the Hyper-V root server and is chosen for our deployments.

Database availability design (DAG) design

What is a DAG? As discussed previously, a DAG is the base component of the HA and site resilience framework built into Microsoft Exchange Server 2010. A DAG is a group of up to 16 Mailbox servers that host a set of databases and provide automatic database-level recovery from failures that affect individual servers or databases.

A DAG is a boundary for mailbox database replication, database and server switchovers and failovers, and for an internal component called Active Manager. Active Manager is an Exchange 2010 component that manages switchovers and failovers that runs on every server in a DAG.

Planning your DAG deployment

From a planning perspective, you should try to minimize the number of DAGs deployed. You should consider going with more than one DAG if you:

- Deploy more than 16 Mailbox servers
- Have active mailbox users in multiple sites (Active/Active site configuration)
- Require separate DAG-level administrative boundaries
- Have Mailbox servers in separate domains (DAG is domain bound)

As required by the customer, each site will have active users. Our design accommodates this requirement by deploying the active/active database distribution model. In the active/active distribution model, active Mailbox database copies are dispersed across both datacenters. Each Mailbox server becomes a primary datacenter for its specific mailbox user population. Passive database copies may exist in alternate data centers for site resiliency. When service for the users of one datacenter fails, the mailboxes containing those users are activated in the other data center.

Three DAGs design

Based on this model, customers can deploy a single DAG that has active mailbox users at each site. However, in the event that the servers in one site temporarily lose connectivity with the other servers in the DAG, the cluster in that site will lose quorum and cease to function correctly. For this reason, this solution deploys three DAGs, with each DAG containing member servers from the primary datacenter that host the primary and secondary database copies. Each DAG also contains servers in one of the alternate datacenters that host the third database copy. The new design has three active/passive DAGs with each datacenter hosting the primary and secondary copies from one DAG, as well as the third copy from another DAG.

The two local database copies are balanced between four Exchange VMs to accommodate an Exchange VM or physical host maintenance/failure scenario.

Table 17, for example, shows that if the Exchange –VM MBX1 on Host 1 fails, its database copies will automatically become activate on MBX2 on Host 2. The same is true if the opposite happens. If the Exchange VM on MBX2 on Host 2 fails, its database copies automatically become active on MBX1 on Host 1.

The same process would occur for a physical server failure. This design provides for the best utilization of the physical hardware and allows for mailbox resiliency within each physical site.

Table 17 provides more details about DAG deployment and the distribution of the database copies among Exchange VMs. In this table, databases are labeled as follows:

- C1 = Active database copy during normal operations
- C2 = Local passive database copy during normal operations
- C3 = Remote passive database copy during normal operations

Table 17. DAG implementation

	Site 1				Site 2				Site 3			
DAG1	MBX1 (Host1)	MBX2 (Host2)	MBX3 (Host3)	MBX4 (Host4)					MBX5 (Host9)	MBX6 (Host10)	MBX7 (Host11)	MBX8 (Host12)
DB1-6	C1	C2							C3			
DB7-12	C2	C1								C3		
DB13-18			C1	C2							C3	
DB19-24			C2	C1								C3
DAG2	MBX13 (Host1)	MBX14 (Host2)	MBX15 (Host3)	MBX16 (Host4)	MBX9 (Host5)	MBX10 (Host6)	MBX11 (Host7)	MBX12 (Host8)				
DB25-30	C3				C1	C2						
DB31-36		C3			C2	C1						
DB37-42			C3				C1	C2				
DB43-48				C3			C2	C1				
DAG3					MBX21 (Host5)	MBX22 (Host6)	MBX23 (Host7)	MBX24 (Host8)	MBX17 (Host9)	MBX18 (Host10)	MBX19 (Host11)	MBX20 (Host12)
DB49-54					C3				C1	C2		
DB55-60						C3			C2	C1		
DB61-66							C3				C1	C2
DB67-72								C3			C2	C1

Placement of the file share witness

The Exchange 2010 DAG feature is built on top of Windows Failover Clustering. A DAG has a quorum when the majority of its members is online and can communicate with the other online members of the DAG. Quorum for the DAG is maintained at the cluster level. It is critical that each DAG member have a consistent view of the DAG's underlying cluster configuration.

The quorum acts as the definitive repository for all configuration information relating to the cluster. The quorum is also used as a tiebreaker to avoid ~~split-brain~~ "split-brain" syndrome. Split-brain syndrome is a condition that occurs when DAG members cannot communicate with each other even though they are up and running. You can prevent split-brain syndrome by ensuring that a majority of the DAG members (and in the case of DAGs with an even number of members, where the file share witness is located on a witness server outside of the DAG) are available and interacting.

The witness server can be any server running the Windows Server operating system, but Microsoft recommends placing it on a highly available server that is handled by the team responsible for managing Exchange. By default, when a DAG is created with an even number of nodes, the file share witness is automatically placed on a Hub Transport server.

When deploying a virtualized Exchange environment, it is common to locate VMs running the Mailbox server role and VMs running the Hub transport role on the same physical Hyper-V root server. Therefore, for all virtualized environments, it is a best practice to locate the file share witness on another highly available server.

Note:

If the witness server you specify is not an Exchange 2010 server, then you must add the Exchange Trusted subsystem universal security group to the local administrators group on the witness server.

In this test environment, file share witness for each DAG is configured on a local domain controller. An alternate file share witness on a domain controller in the remote site is also configured. In a production environment, we do not recommend placing the file share witness on a domain controller due to security implications. You should only use this location as a last resort.

Chapter 4: Storage Design

Overview

This chapter contains the following topics:

Topic	See Page
Exchange Mailbox server storage design with thin provisioning	38
Methodology for sizing Exchange storage	39
Exchange storage design using building block methodology	39
Storage design	47

Exchange Mailbox server storage design with thin provisioning

EMC CLARiiON's thin provisioning feature provides a great benefit for customers deploying large Exchange mailboxes. Physical space is only assigned from a thin pool to the thin LUNs, as needed, using 1 GB increments, up to the logical size specified for each LUN. You can also dynamically expand a thin pool by adding more disks without disruption or requiring downtime. Upon expansion, a thin pool can easily be rebalanced so that the data and workload are wide-striped evenly across the current and newly added disks that make up the pool.

Thin LUN technology works with CLARiiON's metaLUN feature and traditional LUNs to provide powerful, cost-effective, flexible solutions. CLARiiON thin LUNs present more storage to an application than is physically available. Storage administrators are free from the time-consuming administrative work of deciding how to allocate disk drive capacity. Instead, an array-based mapping service builds and maintains all of the storage structures based on a few high-level user inputs. Disk drives are grouped into storage pools that form the basis for provisioning actions. Physical storage is automatically allocated only when writing new data blocks.

With thin provisioning, administrators can provision the minimum amount of required storage to the host while allowing the host to see fully allocated storage. For example, you can provision a 1.2 TB LUN to an Exchange host to satisfy the capacity requirements for 450 users with 2 GB mailbox. After creating a file system on this LUN, and populating mailboxes to 600 MB, only one-third of the LUN (400 GB) is used. At that point, the CLARiiON array intelligence figures out that the remaining 800 GB is not being used and keeps it in reserve. Yet the host continues to think it has access to the remaining 800 GB.

Methodology for sizing Exchange storage

Sizing and configuring storage for use with Microsoft Exchange Server 2010 can be a complicated process, driven by many factors, which vary from organization to organization. Properly configured Exchange storage, combined with a properly sized server and network infrastructure, can guarantee smooth Exchange operations and an excellent user experience. This solution uses the building-block approach to simplify sizing and configurations of storage used with Microsoft Exchange Server 2010. This approach helps exchange storage administrators to deploy large amount of Exchange storage on EMC Unified Storage more easily and efficiently.

Microsoft and EMC provide tools to help you properly size your Exchange Mailbox server. The Exchange 2010 Mailbox server Role Requirements Calculator tool from Microsoft provides CPU and memory guidance's in addition to storage recommendations. EMC's Exchange 2010 storage calculator tool is specifically designed and tailored to provide more in-depth details and accurate recommendations for deploying Exchange on EMC storage. Specifically, the EMC calculator includes thin provisioning features for sizing Exchange storage.

Make sure to consult with a server and storage vendor for additional guidelines during the design and deployment phases. You can download these tools from the following locations:

- For access to the Exchange 2010 Mailbox server Role Requirements Calculator from Microsoft, visit the Microsoft Exchange Team Blog at <http://msexchangeteam.com/archive/2009/11/09/453117.aspx>.
- For the EMC Exchange 2010 storage calculator visit Powerlink at <http://powerlink.emc.com>.

Exchange storage design using building block methodology

What is a building block?

A building block represents the required amount of resources required to support a specific number of Exchange 2010 users on a single VM. You derive the number of required resources from a specific user profile type, mailbox size, and disk requirement. Using the building-block approach removes the guesswork and simplifies the implementation of Exchange VMs.

After designing the initial Exchange Mailbox server VM building block, you can easily reproduce it to support all of the users in your organization that share similar user profile characteristics. By using this approach, Exchange administrators can create their own building blocks based on their company's Exchange environment requirements. This approach is very helpful when a customer expects future growth, as it makes Exchange environment additions easy and straightforward. You can apply this methodology when Exchange is deployed in a either physical or virtual environments and when storage is fully provisioned or thin provisioned (such as in this solution).

EMC's best practices involving the building block approach for an Exchange Server

design has been very successful for many customer implementations. To Create a building-block for a Mailbox server role VM, you need to:

1. Identify user requirements
2. Identify Exchange VM requirements
3. Identify and calculate storage requirements based on both IOPS and capacity
4. Finalize the Exchange VM building block

The following sections detail these four steps.

Step 1. Identify user requirements

Exchange administrators can create building blocks based on the user requirements in their organizations. To obtain user profile information for your existing Exchange environment use Microsoft Exchange Server Profile Analyzer tool available at <http://www.microsoft.com/downloads/details.aspx?familyid=C009C049-9F4C-4519-A389-69C281B2ABDA&displaylang=en>.

Table 18 summarizes key customer requirements for this solution. This information is required to perform the Exchange storage design.

Table 18. Exchange environment requirements

Parameter	Value
Target message profile (messages sent/received/ user/day)	100 messages (0.10 IOPS)
Additional IOPS overhead (per customer requirements)	50%
Target average message size (KB)	75 KB
Outlook mode	100% MAPI
Initial average mailbox size (MB)	Up to 600 MB
Target maximum mailbox size (MB)	2 GB
Total number of mailboxes in the environment	32,400
Total number of users per site	10,800
Number of active users per Mailbox server	2,700
Number of sites	3
Deleted items retention window ("dumpster") (days)	14
Logs protection buffer	5 days
Drive type, speed and capacity	450 GB FC, 15k rpm drive
Database read/write ratio	3:2 in mailbox resiliency configurations

One of the customer requirements for this Exchange 2010 solution is to provide a

two GB mailbox capacity for each of the 32,400 users (10,800 users per site) with the initial mailbox capacity of up to 600 MB. Therefore, our design should incorporate performance and space requirements based on a 2 GB mailbox per user. Make sure to calculate the initial storage capacity requirements so that you purchase only the capacity required at implementation.

Step 2. Identify Exchange VM requirements

Based on the DAG design and the allocation of Exchange Mailbox role VMs per Hyper-V host, four Exchange VMs per site will host 5,400 users with two database copies (2,700 active and 2,700 passive) and the other four VMs will host only 2,700 passive users. Earlier in the design process, we identified CPU and memory requirements for Exchange Mailbox role VMs. Table 19 summarizes these requirements.

Table 19. Exchange VM requirements

VM Role	vCPUs per VM	Memory
Mailbox (5400 Mbx)	4	43 GB
Mailbox (2700 Mbx)	2	24 GB

For this solution, we have eight Exchange Mailbox role VMs per site (24 VMs for a three-site deployment).

Step 3. Identify and calculate storage requirements based on IOPS and capacity

As a best practice for calculating Exchange storage, always calculate both IOPS, and then capacity requirements. The procedures described here show the basic calculations for a targeted user profile. The customer requires that their current infrastructure, which includes an EMC Unified CX4-480 array with 450 GB 15k rpm drives, be incorporated into the design.

Based on the DAG design in this solution, each database supports 450 active mailboxes. Each mailbox server supports six databases or 2,700 active mailboxes on six DB LUNs and six Log LUNs. Therefore, we will use 12 LUN building block supporting increments of 2,700 mailboxes.

Based on this methodology, each Exchange site requires 12 building blocks:

- Four VMs with 5,400 users, with two building blocks each, totaling eight building blocks
- Four VMs with 2,400 users, with one building block each totaling four building blocks

Calculate the mailbox I/O requirements

It is important to understand the amount of database I/O per second (IOPS) consumed by each mailbox user because it is one of the key transactional I/O metrics needed to adequately size your storage. Pure sequential I/O operations are not factored in the IOPS per Mailbox server calculation because storage subsystems can handle sequential I/O much more efficiently than random I/O. These operations include background database maintenance, log transactional I/O, and log replication I/O.

This step describes how to calculate the total IOPS required to support all mailbox

users using the building block methodology.

Note:

To determine the IOPS for different message profiles, refer to the table provided at the following Microsoft TechNet location: [Understanding Database and Log Performance Factors](#)

The total transactional required IOPS = IOPS per mailbox user * number of mailboxes * I/O overhead factor:

$$0.10 * 2700 * 20\% = 324 \text{ IOPS per Exchange VM with 2700 users}$$

Note:

Twenty percent overhead includes log IOPS and BDM IOPS.

In the above procedure, we determined the IOPS requirements to support one building block of 2,700 users. To support 10,800 users per site, we need four of these building blocks. Moreover, to support three database copies per site, we need 12 building blocks. Therefore, the total IOPS required per site is 3,888 (324 IOPS * 12 building blocks) and the all three sites together require 11,664 IOPS.

Calculate the IOPS disk requirements

You calculate the number of disks required to provide the desired user performance, based on the IOPS requirements by using the formula shown below.

(User IOPS x Read Ratio) + Write Penalty (User IOPS x Write Ratio)/IOPS capability of disk type chosen

$$(324 * .6) + 4(324.5 * .4) / 155 = 4.6 \text{ (round-up 5 disks for RAID 5)}$$

Note:

IOPS capacity per disk type can vary depending on the disk type, storage array model, and cache capacity available. Contact your EMC representative to obtain the latest guidelines for disk types and speeds.

Our IOPS calculations concluded that:

- To support 2,700 users with 324 IOPS, five disks are required.
- To support 10,800 users per site with 1,296 IOPS, 20 disks are required.
- To support 32,400 IOPS per site with 3,888, 60 disks are required.

Calculate the mailbox size on disk

It is important to determine the mailbox size on disk before attempting to determine your total storage requirements. A full mailbox with a 2 GB quota requires more than 2 GB of disk space because we have to account for the:

- Prohibit send/receive limit
- Number of messages the user sends/receives per day
- Deleted item retention window (with or without calendar version logging and single item recovery enabled)
- Average database daily variations per mailbox

You can use the Microsoft Mailbox Server Role Requirements Calculator to calculate this number, but we have provided the raw calculations below if you prefer to do

them manually.

Use the following calculations to determine the mailbox size on disk for this solution based on an initial mailbox size of 600 MB and a fully provisioned mailbox quota of 2 GB:

Mailbox Size On Disk = Mailbox Limit + Whitespace + Dumpster

Whitespace = 100 messages / day x 75/1024 MB = 7.3 MB

Dumpster for 600 MB mailbox = (100 messages / day x 75/1024 MB * 14 days) + (600 MB x 0.012) + (600 MB x 0.058) = 144.2 MB

Dumpster for 2 GB mailbox = (100 messages / day x 75/1024 MB * 14 days) + (2048 MB x 0.012) + (2048 MB x 0.058) = 246 MB

Table 20 details the summary of the Mailbox size on disk requirements.

Table 20. Mailbox size on disk summary

Mailbox quota	Dumpster size (2 weeks)	White space	Total size on disk
600 MB	7.3 MB	144.2	752 MB
2 GB	7.3 MB	246 MB	2,301 MB (2.25 GB)

Calculate the capacity requirements and LUN sizes

There are two sets of calculations used for determining the capacity requirements when using thin provisioning:

- Calculations based on the initial capacity requirements. This is necessary to identify the storage requirements to support the initial mailbox capacity. You base the storage capacity purchase on these calculations.
- Calculations based on 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 future.

When leveraging thin provisioning on a EMC Unified Storage, it is a best practice to separate the Exchange logs from the database thin pools. Since log volumes do not have the same growth pattern as the database volumes, it makes sense to separate them. This also provides the flexibility to put log volumes on different disk types or different RAID levels than the database volumes.

Following this best practice, we separate the DBs and logs onto different LUNs. This allows the log volumes to be located on RAID 1/0 LUNs for better performance. Unless the user profile changes significantly in the future, customer do not anticipate any increased capacity for the logs. Therefore the log LUNs are not thin provisioned. More details are provided in Table 14 describes the logs and database capacity requirement calculations.

Database capacity requirements

To determine the actual database size, use the following formula:

Database Size = <Number of Mailboxes> x <Mailbox Size on Disk> x <Database Overhead Growth Factor>

Based on the number of mailboxes, the actual size of the mailboxes, and the database growth overhead factor of 20 percent, the database size is 2,379 GB for initial capacity and 7,290 GB for fully provisioned capacity as shown in the following table.

Table 21. Database capacity requirements

Mailboxes per server	Database size requirements (initial capacity for 600 MB mailbox)	Database size requirements (fully provisioned capacity for 2 GB mailbox)
2,700	2,379 GB (2,700 users * 752 MB + 20%)	7,290 GB (2,700 users * 2.25 + 20%)

Database LUN sizes

To determine the total database LUN size requirements for 2,700 users, use the following:

Database LUN size = <Database Size> + <Content Index Catalog / (1 - Free Space Percentage Requirement)

Note:

The content Index is 10% of the database size.

The following table provides a summary of the calculations. The calculations for database LUN sizes based on the initial mailbox capacity of 600 MB is provided only for references, as they are not being used in this solution. The fully provisioned LUN size based on a 2 GB mailbox size is used instead.

Table 22. Database LUN sizes summary

	Total Database size per server	Content index size (10% of DB size)	Total Database LUN size
Initial capacity for 600 MB mailbox	2379 GB	238 GB (2379 * 0.1)	3,271 GB ((2379 + 238) / .8)
Fully provisioned capacity for 2 GB mailbox	7290 GB	729 GB (7290 * 0.1)	10,024 GB ((7290 + 729) / .8)

To support:

- Twelve building blocks for 10,800 with an initial mailbox quota of 600 MB, we require a 39,252 GB database capacity
- Thirty-six building blocks across three sites with 30,400 users, we require 117,756 GB database capacity
- Two building blocks for 10,800 with fully provisioned mailbox quota of 2 GB, we require 120,288 GB database capacity
- Thirty-six building blocks across three sites with 30,400 users, we require 360,864GB database capacity.

Log capacity requirements

To ensure that the Mailbox server does not sustain any outages because of space allocation issues, make sure to size the transaction logs LUNs to accommodate all of the logs that will be generated during the backup set. If this architecture leverages the mailbox resiliency and single item recovery features as the backup architecture, the log capacity should allocate three times the daily log generation rate in the event that a failed copy is not repaired for four days. (Any failed copy prevents log truncation from occurring.)

A 100-message per day profile mailbox generates 20 transaction logs per day on average, so 2,700 users will generate 54,000 transaction logs each day. With six databases per server and 450 users per database, this means that each database will generate 9,000 logs per day. One percent of the mailboxes move per week and on one day during the weekend. The solution leverages the native data protection features within Exchange and therefore, does not perform backups. The solution is sized to tolerate four days without log truncation.

Calculate the log capacity requirements based on the fully provisioned mailbox size. This way there is no need to make any changes later as log LUNs are not thin provisioned in this solution. Table 23 details summary of log capacity requirements.

Log capacity is calculated using the following formula:

$$\text{Log capacity} = \text{<Log size>} + \text{<Mailbox Move \%>} + \text{<Truncation failure tolerance>}$$

Table 23. Log size requirements

Number of databases per server	Logs per database	Log file size	Daily log size	Move mailbox size / database	Truncation failure tolerance	Log size requirements
6	9000	1 MB	9 GB	10.2 GB (27 × 2.25 GB / 6)	36 GB (4 days × 9 GB)	278 GB (10.2 GB + 36 GB) * 6

Determine the required log LUN size

We determined that 55 GB for log capacity is required to support 450 users per database with 2,700 users per server. Next, we need to calculate the total log LUN size using the following formula. Table 24 details the summary of the calculations.

$$\text{Log LUN size} = \text{<Log capacity >} / (1 - \text{Free Space Percentage Requirement})$$

Table 24. Log LUN size summary

Log size	Number of Log LUNs per server	Free Space Percentage Requirement	Total Log LUN size	Log LUN size per database
278 GB	6	20 %	348 GB ((278 GB) / .8)	58 GB (348 GB / 6)

Six 58 GB log LUNs will be presented to the host with total LUN capacity requirements of 348 GB.

As specified in the DAG design, each site is configured to handle three copies of the

data. This requires 12 building blocks configured across eight VMs. The total Log capacity requirements for one site is 4,176 GB as shown in the following formula.

$$348 * 12 = 4,176 \text{ GB per site and } 12,528 \text{ GB for all 3 sites.}$$

Database disks

We have determined that the 2,700 mailbox building block initially provisioned for a mailbox of 600 MB required a storage capacity of 3271 GB for DB LUNs. The useable capacity per 450 GB spindle in a RAID 5 configuration on the CX4-480 is approximately 402 GB. To determine the number of spindles required use the following formula:

$$\text{Disk Count} = (\text{Total Capacity Required}) / (\text{Useable Capacity} / \text{Spindle w RAID5})$$

$$3271 \text{ GB} / 402 \text{ GB} = 8.1$$

We can accommodate the database capacity requirements with nine disks.

EMC best practices for deploying storage on EMC Unified Storage using thin provisioning are to configure RAID 5 thin pools in multiples of five disks. Therefore, we allocate 10 disks for one building block of 2,700 mailboxes and with extra headroom for future growth.

Log disks

We have also determined that the 2,700-mailbox building block requires a storage capacity of 330 GB for log LUNs. Using two 450 GB drives in a RAID 1/0 configuration on a CX4-480 provides 402 GB of usable storage capacity. Therefore, the proposed two-disk configuration meets the log capacity requirements of the 2,700-mailbox building block.

There are four building blocks supporting active mailbox databases on the primary datacenter mailbox servers. There are four building blocks supporting passive mailbox databases on the primary datacenter mailbox servers. There are four building blocks supporting passive mailbox databases on the secondary datacenter mailbox servers. Therefore, we require 24 spindles to support log LUN requirements in a single site.

Step 4. Finalize the Exchange VM building-block

By performing IOPS and capacity calculations, we have determined that the total capacity requirements supersede the IOPS requirements.

Exchange VMs designed to handle only passive copies (2,700 users) require one building block of 10 spindles from a thin pool. Exchange VMs designed to handle either the active and passive copies, or 5,400 users (2,700 active/2,700 passive), require two building blocks with 20 spindles from a thin pool.

Storage design summary

Building block summary

Table 25 summarizes the configuration for a building block of 2,700 users. A Mailbox server handling two copies requires two of these building blocks, for a total of 5,400 users per server.

Table 25. Building-block summary (based on 600 MB initial capacity)

Users per Exchange VM	Disks per VM	CPU's per VM	Memory per VM
2,700	12 (10 for DBs, 2 for logs)	2	24 GB

Capacity requirements summary

Our capacity calculations concluded that using the building block methodology for calculating Exchange mailbox storage provides an easy and effective method. Adding more users with the same profile involves allocating the same amount of storage, memory, and CPU resources to achieve the necessary performance. This flexible design offers customers the capability of keeping pace with an increasing user population. You can easily add users that share the same profile to the environment, as shown in Table 26.

Table 26. Capacity requirements summary(based on 600 MB initial mailbox capacity)

To support...	Requires...
2,700 users	12 disks (10 for DBs, 2 for logs)
5,400 users	24 disks (20 for DBs, 4 for logs)
10,800 users	56 disks (40 for DBs and 16 for logs)
32,400 users	144 disks (120 for DBs and 24 for logs)

Total storage requirements summary

Table 27 summarizes the required storage and savings that you can achieve by using EMC thin provisioning in this solution.

Table 27. Thin provisioning storage savings

Requirements	10,800 users per site (3 copies and Initial mailbox size 600 MB)	10,800 users per site (3 copies and thin provisioned 2 GB mailbox)	Initial savings
Database space requirements	39,252 GB	120,288 GB	81,036GB (243,108 for 3 sites)
Log space requirements	12,528 GB	12,528 GB	N/A

In summary, CLARiiON Virtual Provisioning features allow customers to provide a total initial savings of over 81 TB of storage across three sites. Customers can purchase and deploy their storage at later time and at a much lower cost as user mailboxes continue to grow beyond the 600 MB initial quota. Table 27 provides a summary of the Exchange server configurations in this solution.

Determine the thin pool strategy

Based on the DAG design and storage requirements identified earlier, we can now determine how many databases to configure for each Exchange Mailbox role VM. As a best practice for thin provisioning Exchange storage on CLARiiON, you should configure the logs on fully provisioned storage using the traditional FLARE® LUNs for better performance.

Now that we have determined the number of spindles required to support the IOPS and capacity requirements of the building block, we need to determine the best way to provision LUNs on the array for that building block when using virtual or thin provisioning. There are three main models you can when designing thin pools for use with Exchange:

- **A storage pool for each Exchange Mailbox server** provides for a more granular (building block) design that is easier for troubleshooting and analysis. This configuration provides the best performance and prevents any disk contention issues that can arise during activities such as seeding/reseeding, backup, and online maintenance (BDM).
- **One storage pool for each Exchange database copy** is a simple method for designing thin pools and provides for best space utilization. This method also ensures that the database copy spindle isolation is maintained if customers must provision additional storage to accommodate growth. However, this method may make it more challenging to perform troubleshooting and performance analysis.
- **One storage pool for all Exchange databases and logs** is another simple method for designing thin pools that provides the best space utilization. However, EMC does not recommend a single thin pool when multiple copies of the same database are located on the same physical array. This creates disk contention during and seeding/ re-seeding operations and during online maintenance (BDM). EMC does not recommend this approach.

In this solution, we separated each database copy in its own thin pool and created three pools of 40 disks on each CX4-480 storage array. Each pool provides support for one copy and 10,800 users. With this configuration, you have more flexibility when adding new users and provisioning new storage when necessary.

The first thin pool contains a 2,700-mailbox building block from each of the four primary datacenter mailbox servers at the site. Previously, we determined that 10 spindles were required to support the IOPS and capacity requirements of the building block. Therefore, the first thin pool supporting 10,800 active mailboxes requires 40 spindles.

The second thin pool also contains a 2,700-mailbox building block from each of the four primary datacenter mailbox servers at the site. The second thin pool supporting 10,800 passive mailboxes requires 40 spindles.

The third thin pool also contains a 2,700-mailbox building block from each of the four secondary datacenter mailbox servers at the site (i.e., the servers from an alternate

DAG that are supporting the site resilient database copies). The third thin pool supporting 10,800 passive mailboxes requires 40 spindles.

Figure 4 shows how the solution distributes the databases from a storage pool for each Exchange Mailbox server VM and Hyper-V host at one of the sites. Mailbox servers from DAG1 hosting 5,400 users are designed with 12 databases, where six of them (2,700 users) are active databases from one storage pool and the other six (2,700 users) are passive from a second storage pool. Mailbox server members from DAG2 hosting only passive copies are designed with only six databases from the third storage pool.

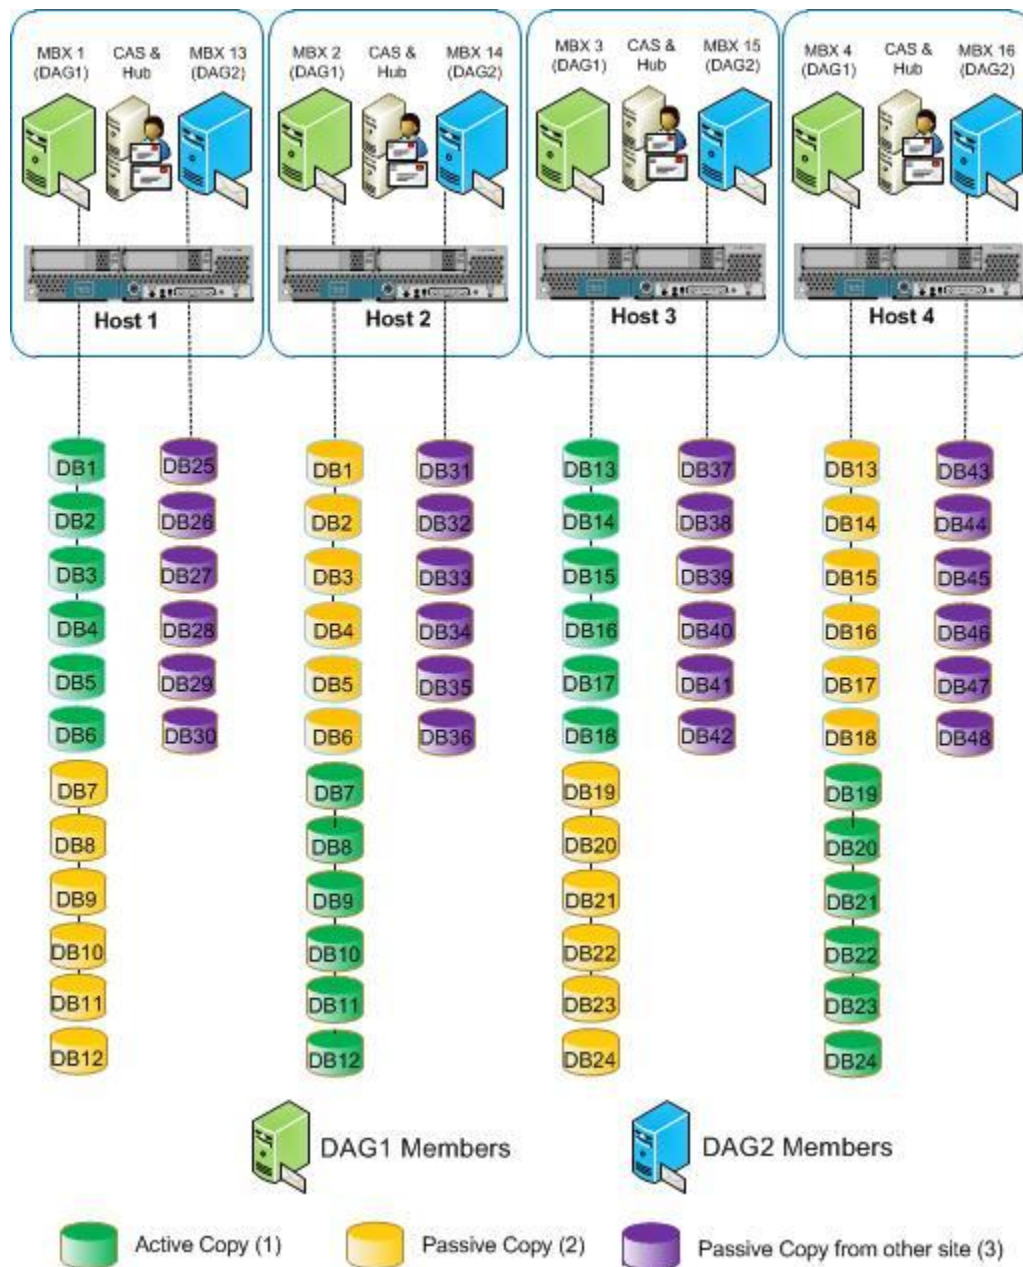


Figure 4. Database distribution per Exchange Mailbox role VMs

Mailbox server configuration summary

In summary, designing Exchange storage for this solution using CLARiiON Virtual Provisioning features allow customers to provide a total initial savings of over 243 TB of storage across three sites. Customers can purchase and deploy their storage as needed and at a much lower cost as user mailboxes continue to grow beyond the 600 MB initial quota.

Table 28 provides a summary of the Exchange Server configurations in this solution.

Table 28. Exchange Server configurations for this solution

Components	Mailbox servers (2,700 users)	Mailbox servers (5,400 users)	HUB/CAS servers
vCPU	2	4	3
Memory	24	43	8
Physical Disks	12 (10 DB, 2 Logs)	24 (20 DB, 4 Logs)	1
LUNs	13 LUNs 1 OS – VHD 12 physical pass-through (6 DBs, 6 logs)	25 LUNs 1 OS – VHD 24 physical pass-through (12 DBs, 12 logs)	2 LUNs 1 OS – VHD 1 physical pass-through for queues transport database

Chapter 5: LAN and SAN Architecture

Overview

This chapter describes the solutions network architecture and contains the following topics:

Topic	See Page
UCS server architecture	51
Fabric Interconnects	51
SAN zoning and LAN connectivity	52
Network load balancing and namespace planning	54

UCS server architecture

Each UCS server is a half-width blade with a mezzanine-connected converged network adapter (CNA) for up to 20 Gbps of I/O throughput. The CNA provides dual-port connectivity to the server chassis midplane and presents the host operating system with two 10 Gigabit Ethernet adapters and two QLogic 4-Gbps FC HBAs.

Each UCS chassis includes two fabric extenders. A fabric extender has four 10-Gigabit Ethernet, a FCoE-capable, several Small Form-Factor Pluggable Plus (SFP+) ports that connect the blade chassis to the Fabric Interconnect, and eight 10-Gigabit ports connected through the midplane to each half-width slot in the chassis. The two fabric extenders provide redundancy and up to 80 Gbps of I/O to the chassis.

Fabric Interconnects

Fabric Interconnects use a cut-through architecture that supports low-latency, line-rate 10 Gigabit Ethernet on all ports, independent of packet size and enabled services. The product provides low-latency, lossless 10 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The Fabric Interconnect supports multiple traffic classes over a lossless Ethernet fabric from the blade through the interconnect.

The FCoE-optimized server design enables network interface cards (NICs), host bus adapters (HBAs), and cables. The server also switches consolidation of the LAN and SAN traffic onto a single unified fabric, thereby reducing expenses associated with multiple parallel networks, different types of adapter cards, switching infrastructure, and cabling within racks. FC expansion modules in the fabric interconnect support direct connections from the UCS servers to native FC SANs.

Each Fabric Interconnect has two 4-Gbps ports connected to one of the two MDS

9134 switch. Each MDS 9134 switch presents an independent and SAN fabric. The CX4-480 storage array controllers have two 4-Gbps ports connected to each of the two MDS 9134 switches.

The fabric interconnects are configured in End Host mode. This default configuration option enables Fabric failover. Fabric failover provides redundant connectivity for Ethernet adapters in the Cisco B200 blades without requiring NIC teaming software. The port redundancy is implemented in the hardware. Both 10 GbE adapters operating in active/active mode provide failover capability for the partner port.

SAN zoning and LAN connectivity

Soft zoning is implemented on the FC switches for SAN traffic isolation. Zone groups are configured to isolate storage I/O between each initiator port on the host and two target ports; each target port on a different storage processor. This SAN configuration provides a redundant SAN fabric with no single point of failure.

LAN connectivity is facilitated by two Nexus 5010 10 GbE converged fabric switches. Each fabric interconnect has a 10 GbE connection to each of the two Nexus 5010 switches. The Nexus 5010 switches are linked together with two 10 GbE connections that facilitates redundant vPC (virtual port channel) interconnect links. A vPC port group that is configured for each pair of Fabric Interconnect connections provides LAN redundancy.

WAN for database replication

The estimated Exchange DAG replication bandwidth requirement between two sites is 155 Mbps with a maximum latency of 50 msec. These values account for the transaction log and content index replication and include 30 percent transaction log compression. The transaction log compression is implemented as part of the Exchange DAG replication. In our solution, each physical site was defined as an AD site. All replication traffic between local database copies (C1 and C2) within an AD site is *intrasite* traffic. This type of traffic is not encrypted or compressed. All replication traffic between local and remote copy (C1 and C3) between AD sites is *intersite* traffic and is encrypted and compressed by default. The values are estimated using the Exchange 2010 Mailbox server Role Requirements Calculator. The RPO is another requirement that we factored into the calculation.

An OC-3 WAN link provides 155 Mbps of bandwidth and meets our requirements. In order to avoid a single point of failure, dual OC-3 links are provisioned for redundancy.

Figure 5 shows the replication-related parameters entered into the Mailbox Role calculator to obtain the WAN requirements for database replication.

Log Replication Configuration	
Hours in the Day	Logs Generated / Hour Percentage
1	2.00%
2	2.00%
3	2.00%
4	2.00%
5	2.00%
6	3.00%
7	3.00%
8	9.00%
9	9.00%
10	6.00%
11	6.00%
12	4.00%
13	8.00%
14	6.00%
15	5.00%
16	5.00%
17	5.00%
18	5.00%
19	4.00%
20	3.00%
21	3.00%
22	2.00%
23	2.00%
24	2.00%
Total	100.00%
Peak Log & Content Index Replication Throughput Requirements	
Peak Log & Content Index Throughput Required / Database	6.48 Mbps
Peak Log & Content Index Throughput Required Between Datacenters / DAG	155.44 Mbps
Total Peak Log & Content Index Replication Throughput Required / Environment	155.44 Mbps
Chosen Network Link Suitability	
Network Link	OC-3 (155 Mbps)
Network Link Latency	50 ms
Is Network Link Acceptable for Peak Replication?	Yes
Is Network Link Acceptable for RPO Replication?	Yes
RPO Log & Content Index Replication Throughput Requirements	
Recovery Point Objective	Near 0 Hours
RPO Log & Content Index Throughput Required / Database	6.48 Mbps
RPO Log & Content Index Throughput Required Between Datacenters / DAG	155.44 Mbps
Total RPO Log & Content Index Replication Throughput Required / Environment	155.44 Mbps

Figure 5. Bandwidth requirements from Microsoft's mailbox role calculator

Network load balancing and namespace planning

Overview

Compared to previous Exchange Server releases, some architectural changes in Exchange 2010 have resulted in network load balancing that is becoming increasingly important, both for large-scale and small-scale deployments.

Exchange RPC Client Access and Address Book services

The addition of the RPC Client Access Service and the Exchange Address Book Service improve the user experience during Mailbox role failovers by moving the connection endpoints for Outlook (and other MAPI clients) to the CAS role rather than the Mailbox role.

In previous Exchange versions, Outlook connected directly to the Mailbox server responsible for the data being accessed, and directory connections were either proxied using the Mailbox role, or they were referred directly to a particular Active Directory Global Catalog (GC). Now that these connections are handled by the CAS role, you must load balance the Outlook connections (both internal and external) across the array of CAS servers in the deployment. In Exchange Server 2010, this is known as a CAS array.

Figure 6 below shows how the solution team configured the CAS arrays.

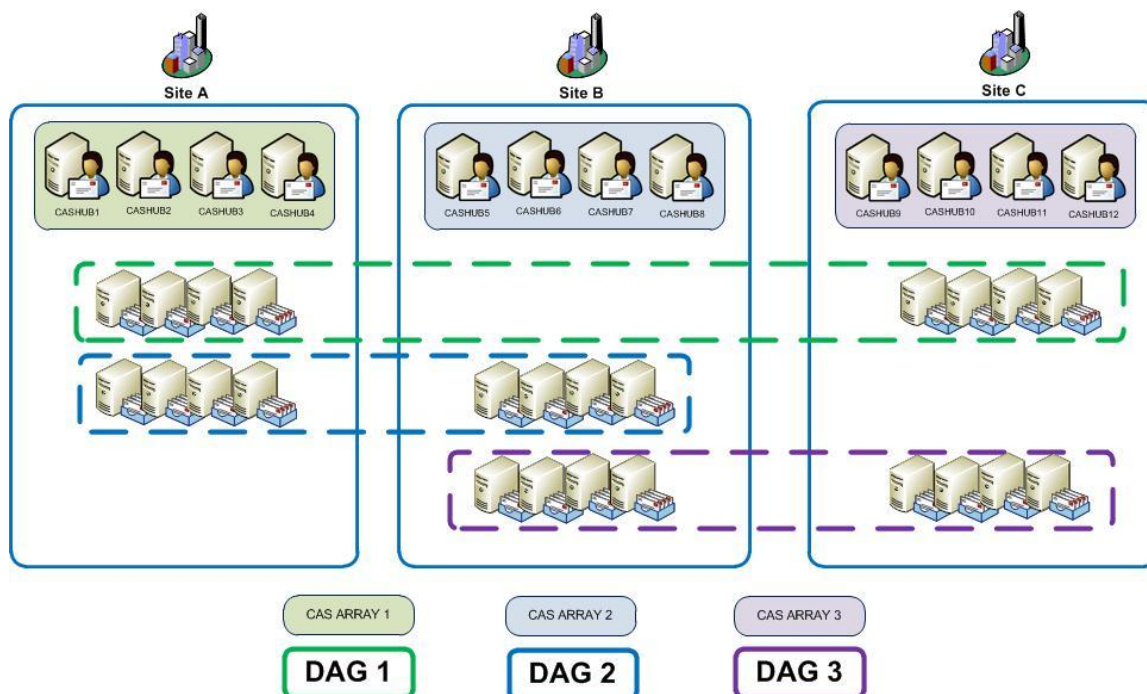


Figure 6. CAS array configuration

Regional namespace model

Because we have active mailbox users in each of the three geographic locations and each of those locations have their own set of Client Access servers, a regional namespace model was deployed. Each location is a separate active directory site with its own DNS namespace (for example, mail.site1.domain.com, mail.site2.domain.com, and mail.site3.domain.com) and a Client Access array defined for that namespace.

The main advantage of this model compared to a multi-site single namespace model is reduced proxying because the majority of mailbox users connect to a Client Access server in the same Active Directory site as their Mailbox server. This improves end-user experience and performance.

Note:

In a single namespace model, it is possible to connect to an out of site CAS server.

Network traffic

Network traffic to each Client Access server array is load balanced by two ACE 4710 units that are configured for redundancy in each site. Each ACE 4710 unit has 4 1-GbE connections to one of two Nexus 5010s in each site. These 4 1-GbE ports are in an individual virtual port channel group.

Monitoring the Outlook client network ports

Since all users are connecting with an Outlook client from the internal network, we are primarily concerned with load balancing TCP socket-oriented traffic. We need to ensure that that traffic maintains client IP-based persistence. Outlook initiates a connection to the RPC Client Access Service and Exchange Address Book Service using the RPC endpoint mapper on port 135.

We set the RPC Client Access server to use static port 60000 and the Exchange Address Book Service to use static port 60001. This port handles connections for both the Address Book Referral (RFR) interface and the Name Service Provider Interface (NSPI). If we do not set static ports, a random port is used. This means that a large range of destination ports may need to be configured for load balancing without the ability to specifically target traffic for these services based on port number.

The ACE 4710s are configured for layer four load balancing by monitoring TCP port 6020 on each Client Access server network interface. TCP port 6020 is a port used for MAPI network traffic. This enables the load balancer to divide the load equally among the available Client Access servers that are members of an Exchange Client Access server array. If a CAS server stops responding on TCP port 6020, the ACE 4710 reconnects Outlook MAPI clients to other accessible Client Access servers in the Client Access server array.

Outlook also uses some HTTP based services including Auto discover, Exchange Web Services (EWS), and Offline Address Book (OAB). These protocols are load balanced. Figure 7 provides a detailed view of the physical infrastructure configuration for this solution.

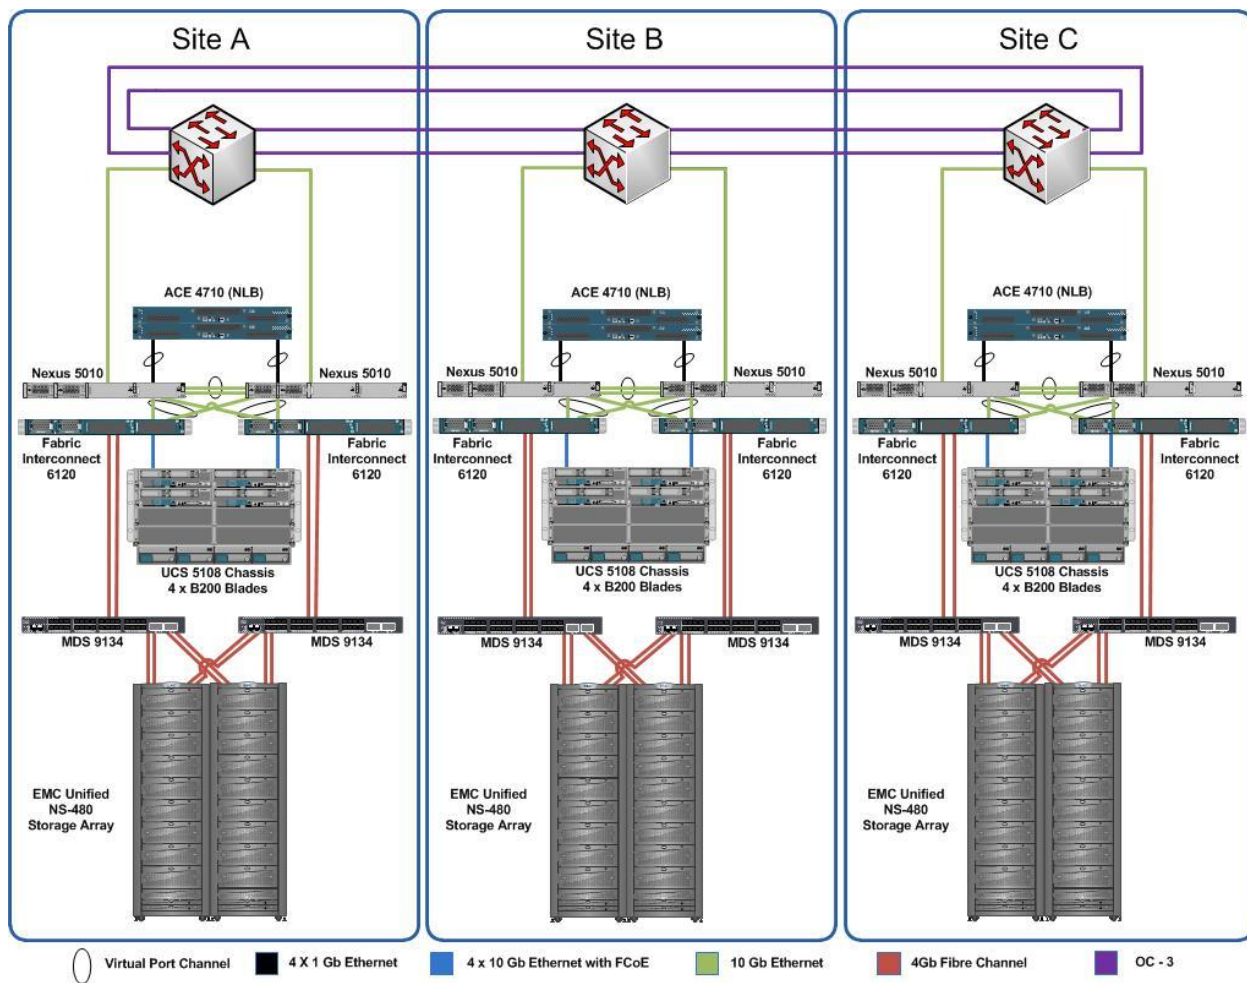


Figure 7. Physical architecture for this solution

Chapter 6: Best Practices Planning

Overview

This chapter contains the following topics:

Topic	See Page
Exchange 2010 best practices	57
Optimizing SAN best practices	57
Mailbox server optimization for EMC storage	58
Unified Computing System best practices and optimization	59

Exchange 2010 best practices

In comparison to earlier versions of Microsoft Exchange, Exchange Server 2010 has made significant improvements in the areas of I/O and storage. For example, there have been many changes to the core schema as well as to Exchange's extensible storage engine (ESE) to reduce the I/O usage profile. Due to this I/O reduction, Exchange 2010 now supports more drive types such as SATA and SAS disks, as well as FC and EFD drives.

For information on Exchange 2010 Mailbox server design best practices, go to <http://technet.microsoft.com/en-us/library/dd346703.aspx>. In addition to Microsoft's recommendations, EMC recommends the best practices described in this section when planning a CLARiiON storage implementation for optimal performance results with Exchange 2010.

Optimizing SAN best practices

Although storage area networks (SANs) provide excellent storage architectures for Exchange implementations, it is important that you optimize your SAN for reliability and performance. The following best practices are important to consider when implementing Exchange in a SAN environment.

To optimize a SAN for reliability, you should:

- Configure redundant controllers, SAN switches, and use RAID.
- Use redundant HBAs connected to different fabrics.

To optimize a SAN for performance, you should:

- Install EMC PowerPath on the physical hypervisor hosts for optimal path management and maximum I/O performance. For more information on installing and configuring the PowerPath application, go to:
-

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

- Dedicate physical spindles within your SAN to Exchange databases to isolate the Microsoft Exchange server database workload from other I/O-intensive applications or workloads. This ensures the highest level of performance for Exchange and simplifies troubleshooting in the event of disk-related issues.
- Plan for performance even in a failover situation. Balance LUNs across the array storage processors to take advantage of CLARiiON's performance and high-availability features
- Plan so that expected peak utilization does not exceed 80 percent saturation of the system.
- Configure the storage to support the IOPS value that you calculated from instructions earlier in this white paper. Always size the Exchange environment for IOPS, then the capacity.

After calculating the IOPS requirements, always apply a 20 percent I/O overhead factor to your calculations to account for additional IOPS, such as logs, log replication, and background database maintenance (BDM) that are not included in the IOPS per user profile. For more information on this topic, see Microsoft's TechNet article "Understanding Database and Log Performance Factors" at <http://technet.microsoft.com/en-us/library/ee832791.aspx>. You should:

- Verify that your SAN switch can support the IOPS requirements, even in a failover situation. The SAN switch has to process the incoming I/O request and forward it to the appropriate port. This therefore limits the amount of I/O that can be handled.
- Verify that the Host Bus Adapter (HBA) installed in the server can support your IOPS requirements, even in a failover situation. To avoid throttling, ensure that the queue depth is set according to EMC recommendations.

Mailbox server optimization for EMC storage

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

- Do not perform partition alignment on Microsoft Windows Server 2008 as partitions are automatically aligned to a 1 MB offset. (Exchange Server 2010 requires Windows Server 2008 or Windows 2008 R2.)
- When formatting a new NTFS volumes for an Exchange database and logs, you should set the allocation unit size (ALU) to 64 KB using the drop down list in Disk Manager or through the CLI using the **diskpart** command.
- Microsoft recommends a maximum database size of 200 GB in environments where DAG is not being used. When using DAG with a minimum of two RAID protected database copies, the maximum database size can be up to two TB. Consider backup (if applicable) and restore times when calculating the database size.
- Enable BDM on large databases (greater than 500 GB).

For more information on EMC solutions for Microsoft Exchange Server, visit the EMC website at: <http://www.emc.com/solutions/application->

environment/microsoft/solutions-for-microsoft-exchange-unified-communications.htm

Unified Computing System best practices and optimization

The following list outlines the UCS best practices for Exchange server deployments:

- Implement redundant components to avoid downtime in case of hardware failure or component failure. The optional redundancy is available for the following components:
 - Fabric interconnect
 - Fabric interconnect power supply
 - Fabric extender
 - Chassis power supply
 - Configure dual HBA adapter ports and PowerPath for SAN redundancy. Configure one HBA port for Fabric A, and another HBA for Fabric B.
 - Configure the Fabric Interconnects for end-host-mode. This option enables fabric failover for LAN adapters.
 - Configure dual network adapters for LAN access redundancy. Configure one adapter to use fabric A and another adapter to use fabric B. Enable both adapters for fabric failover. You do not need to deploy network adapter teaming software in this case.
 - Configure redundant network links for the fabric interconnect uplink Ethernet ports using virtual port channel (vPC) or similar technologies.
 - Separate DAG replication traffic and all other network traffic on different adapters. Configure one adapter to carry the DAG replication traffic and another adapter to carry the all other network traffic, including MAPI traffic. The DAG replication traffic should use a different TCP/IP subnet than all other network traffic.
 - Place the Active Directory servers on a separate IP subnet from the Exchange servers.
 - Use UCS Manager service profile pools to allocate WWPN, WWNN, MAC addresses, and server UUIDs. This option enables these attributes to move from one server blade to another server blade with the service profile. One benefit of this approach is a case where a blade server fails and needs to be replaced with new hardware. You can replace the blade server in this case without reconfiguring the SAN zoning, LUN masking, VLAN assignments, or static DHCP reservations.
-

Chapter 7: Solution Validation

Overview

Introduction

This chapter describes the approach and methodology used to validate this solution, which involves both functional and performance tests.

The performance tests include:

- Storage performance validation using Jetstress
- Database seeding performance
- Server and Exchange environment (end-to-end) performance validation using Loadgen

The functional tests included a site (datacenter) failover/failback validation.

Contents

This chapter contains the following topics:

Topic	See Page
Methodology and tools	61
Exchange storage validation with Jetstress	64
Database seeding performance	67
Environment validation with Loadgen	68
Validating primary datacenter service restoration	78

Methodology and tools

Overview

You can use a variety of tools to measure the performance of Exchange 2010, including *Jetstress* and *Load Generator* (LoadGen). The Windows Server 2008 operating system also includes some general performance tools including [Windows Performance Monitor](#). You can perform EMC Unified Storage analysis using EMC Unisphere Manager.

In addition to these tools, you should analyze your current user loads to establish a minimum server requirements baseline. Understanding how your users use the system is one of your biggest challenges. The Exchange Server Profile Analyzer can help provide useful data when analyzing your current user loads. After you determine your hardware requirements, you should conduct a pilot test to make sure performance levels are acceptable.

For more information, see [Tools for Performance and Scalability Evaluation](#) available on Microsoft's TechNet website.

Jetstress 2010

The best tool for validating the Exchange storage design is Jetstress. Jetstress simulates Exchange I/O at the database level by interacting directly with the database technology of the ESE, also known as the Jet on which Exchange is built. You can configure Jetstress to test the maximum I/O throughput available to your disk subsystem within the required performance constraints of Exchange. You can also configure it to accept a user count, profile, and I/O per second per user to validate that the disk subsystem is capable of maintaining an acceptable performance level with that profile. We strongly recommend that you use Jetstress to validate storage reliability and performance prior to the deploying your Exchange servers in to a production environment.

You can download Jetstress from [Microsoft Exchange Server Jetstress 2010 \(64 bit\)](#) at <http://go.microsoft.com/fwlink/?LinkId=178616>. The Jetstress documentation describes how to configure and execute an I/O validation or evaluation on your server hardware.

While the Jetstress tool tests the performance of the Exchange storage subsystem before placing it in the production environment, it does not test the impact of the server CPU and memory configuration of MAPI user activity. Use the Microsoft Loadgen tool for this purpose.

Loadgen

You use Exchange Load Generator (Loadgen) to perform a full end-to-end assessment of the Exchange 2010 environment. You can use Loadgen to perform pre-deployment validation and stress testing tasks that introduce various workload types into a test (non-production) Exchange messaging system. This test simulates the delivery of multiple MAPI, Outlook Web access, IMAP, POP, and SMTP client messaging requests to an Exchange server.

Important!

You should use Loadgen only in a test lab configuration and in non-production Exchange environments. For more information on Loadgen go to the following website:

<http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=cf464be7-7e52-48cd-b852-ccfc915b29ef>

EMC Unisphere

EMC Unisphere provides a flexible, integrated experience for managing existing CLARiiON storage systems and next-generation EMC unified storage offerings. This new approach to mid-tier storage management fosters simplicity, flexibility, and automation. Unisphere's unprecedented ease-of-use is reflected in intuitive task-based controls, customizable dashboards, and single-click access to real time support tools and online customer communities.

New Unisphere features include:

- Task-based navigation and controls that offer an intuitive, context-based approach to configuring storage, creating replicas, monitoring the environment, managing host connections, and accessing the Unisphere support ecosystem.
- A self-service Unisphere support ecosystem is accessible with one click from Unisphere. This feature provides users with quick access to real time support tools, including live chat support, software downloads, product documentation, best practices, FAQs, online communities, spare orders, and service requests.
- Customizable dashboard views and reporting capabilities that enable at-a-glance management. Unisphere automatically presents users with valuable information based on how they manage their storage.
- A single sign-on and integrated experience for managing CLARiiON and Celerra® platforms.

Figure 8 provides an example of the Unisphere Summary page that gives administrators a wealth of detailed information on connected storage systems, from LUN pool and tiering summaries to physical capacity and RAID group information.

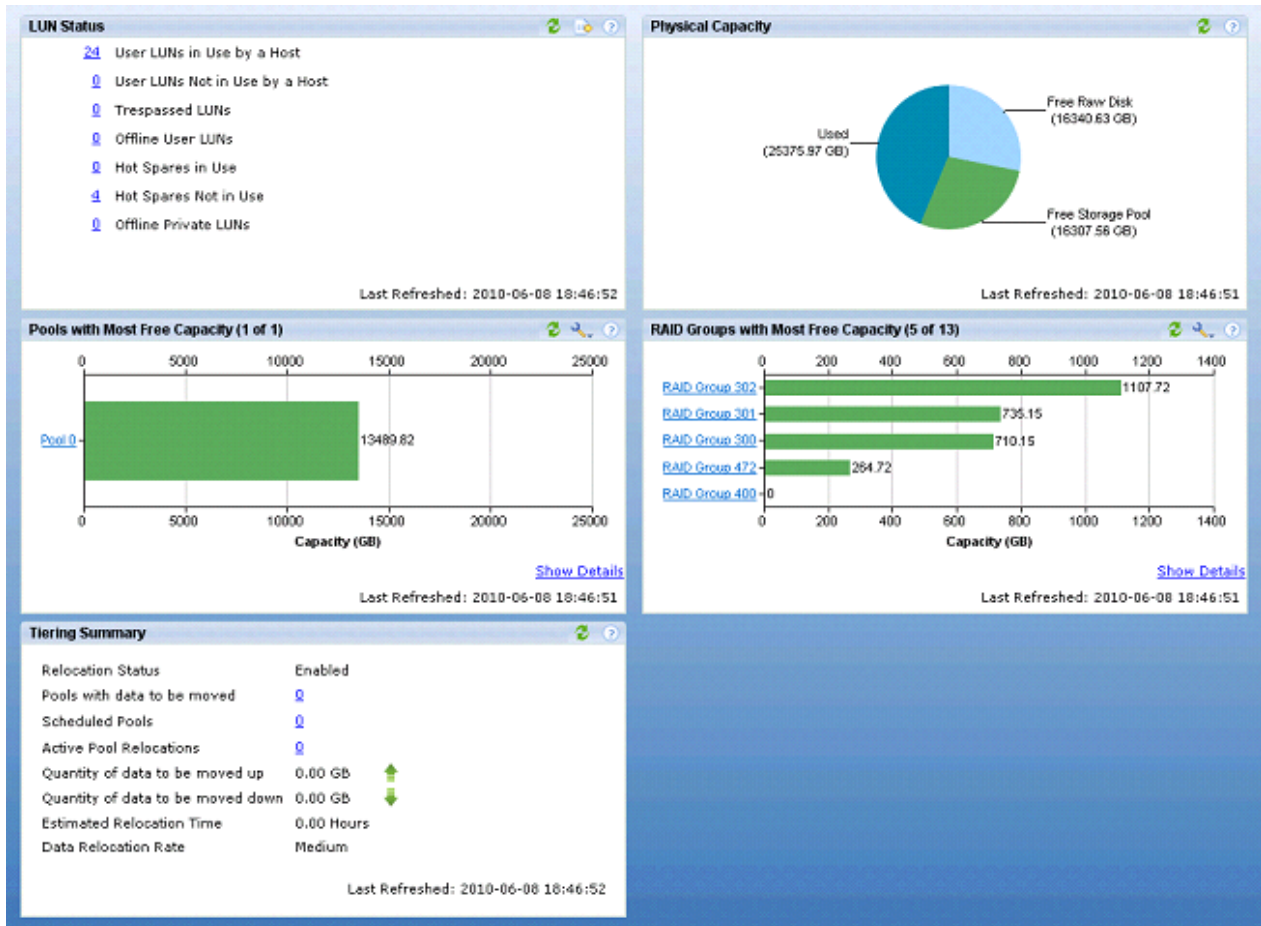


Figure 8. Unisphere summary page

Exchange storage validation with Jetstress

Overview

Before implementing a storage solution in a production environment, it is important to verify that the Exchange storage is sized and configured properly. This section describes the approach and methodology that the solution team used to validate the storage design. The testing performed is similar to the one required for Microsoft's Exchange Solution Reviewed Program (ESRP) program, which is designed for storage vendors like EMC to submit their Exchange solutions.

The Microsoft's ESRP program provides a common storage-testing framework for vendors to provide information on its storage solutions for Microsoft Exchange Server software. The ESRP program is not designed to be a benchmarking program; tests are not designed to measure the maximum throughput for a given solution. ESRP testing focuses on producing recommendations from vendors for supported Exchange configurations. For more details on the Microsoft ESRP storage program visit:

<http://technet.microsoft.com/en-us/exchange/ff182054.aspx>

Currently, Microsoft does not accept ESRP submissions from storage vendors if Exchange Mailbox servers are configured in a virtualized environment (VMware® or Microsoft Hyper-V). That is why EMC finds it important to provide customers with vital data points and performance information to help deploy or transition Exchange from a physical to a virtual environment.

Test results listed in this white paper are provided according to ESRP guidelines for easy comprehension and comparison to other ESRPs published on Microsoft's website (including ESRP submissions from EMC and other storage providers).

Test configuration

To validate all disks configured for Exchange on a single CX4-480 array, the solution team ran Jetstress against all disks configured for three database copies simulating a 32,400-user load. Although our DAG design has no conditions in which all three copies are active, you should perform the testing against all disks. This is necessary to ensure that all Exchange storage is configured properly to sustain switchover and failover conditions. We have performance and shows results for a Jetstress test against all disks configured for three database copies, or 32,400 users across eight Mailbox servers on a single CX4-480 storage array.

Jetstress was run simultaneously from all eight Exchange Mailbox servers in one site at the same time. Four Mailbox servers were configured to each simulate the load for 5,400 users at .18 IOPS per user and the other four servers to each simulate the load of 2,700 users at 0.18 IOPS per user.

Jetstress test results and CX4-480 performance

Figure 9 and Table 29 display the test results for a single CX4-480 array performance for a Jetstress test against all disks configured for three database copies, or 32,400 users across eight Mailbox servers.

Testing against all disks on a single storage frame shows that CX4-480 achieved 8,120 Exchange 2010 transactional IOPS across eight Exchange VMs. This is 4,232 IOPS over the designed target baseline of 3,888 IOPS for a single site. This

additional headroom provides a nice buffer and insurance against any unexpected DAG failovers, large mailbox moves, I/O spikes, peak loads, and potential malware attacks that may have otherwise taken the server down. Disk latencies were all within the acceptable parameters according to Microsoft's best practices for Exchange 2010 performance.

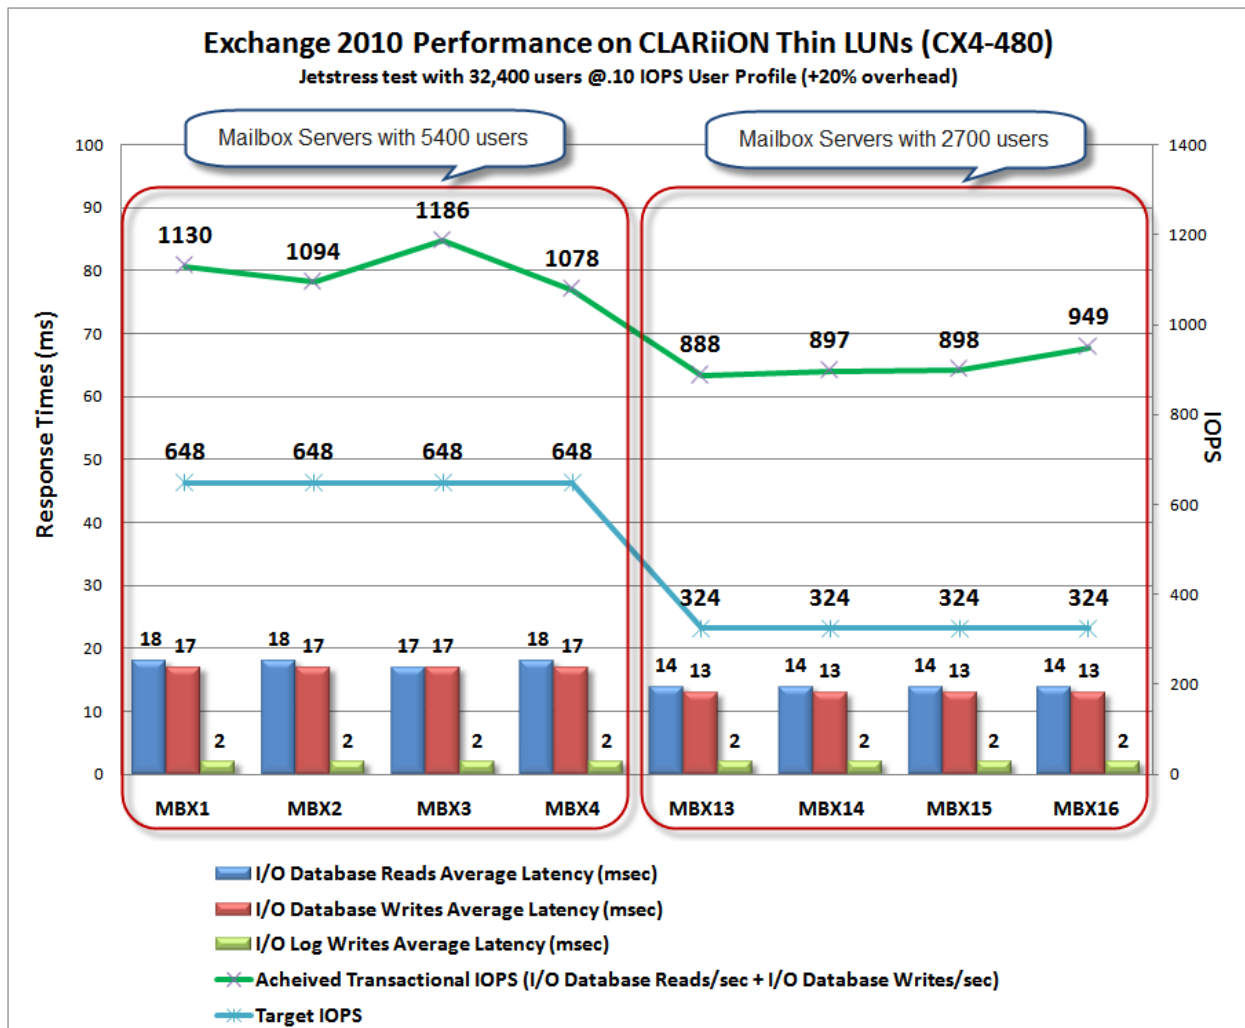


Figure 9. Exchange 2010 performance on CX4-480 with thin LUNs.

**Jetstress
performance
metrics and
test results
summary**

Table 29 shows the aggregate performance across all servers, which is the sum of all I/Os across the solution's servers and the average latency across all of the solution's servers.

Table 29. Jetstress test results summary

	Target values	4 Mailbox servers with 2,700 users	4 Mailbox servers with 5,400 users	Total
Database I/O				
Achieved transactional IOPS (I/O database reads/sec + I/O database writes/sec)	For 2700 users - 1296 IOPS For 5400 users - 2592 IOPS	3632 IOPS	4488 IOPS	8120
I/O database reads/sec	N/A	2193	2729	4922
I/O database writes/sec	N/A	1439	1759	3198
I/O database reads average latency (msec)	<20 ms	14	18	N/A
I/O database writes average latency (msec)	<20 ms	13	17	N/A
Transaction log I/O				
I/O log writes/sec	N/A	1242	1560	2802
I/O log reads average latency (msec)	<10 ms	2	2	2

Database seeding performance

Overview

Database mobility is a new architecture in Microsoft Exchange Server 2010 that removes the concept of storage groups and uncouples an Exchange 2010 Mailbox database from a Mailbox server. Because Microsoft removed storage groups from Exchange 2010, continuous replication now operates at the database level. In Exchange 2010, transaction logs are replicated to one or more Mailbox servers and replayed into one or more copies of a mailbox database stored on those servers.

For a variety of reasons, such as performing planned maintenance, it may be necessary to suspend and resume continuous replication activity for a database copy. In addition, some administrative tasks, such as seeding, require that you suspend a database copy first. Microsoft recommends that you suspend all replication when you change the path for the database or its log files. You can suspend and resume database copy activity by using the Exchange Management Console, or by running the **Suspend-DatabaseCopy** and **Resume-DatabaseCopy** cmdlets in the Shell.

For detailed steps that explain how to suspend or resume continuous replication activity for a database copy, visit Microsoft's TechNet Website at [Suspend or Resume a Mailbox Database Copy](#).

Note:

Log truncation does not occur on the active mailbox database copy when one or more passive copies are suspended. If your planned maintenance activities are going to take an extended period of time (for example, several days), you may have considerable log file buildup. That is why it is critical to size your log LUNs to accommodate for these activities. To prevent the log drive from filling up with transaction logs, you can remove the affected passive database copy instead of suspending it. When the planned maintenance is completed, you can re-add the passive database copy.

Database seeding considerations

While the database is being seeded, the backup cannot occur. Careful planning is required to seed all databases in the entire environment.

In this solution, 72 database copies are configured. 24 of these database copies are active in each of the three sites and other 48 are passive.

We have seeded the databases in the following order at each site:

Local Primary → Local Passive → Remote Passive

The seeding source for the remote passive copy was the local passive copy.

Performing a seed operation

When performing a seed operation, you can choose to seed the mailbox database copy, the content index catalog for the mailbox database copy, or both the database copy and the content index catalog copy. The default behavior of the Update Database Copy wizard and the [Update-MailboxDatabaseCopy](#) cmdlet is to seed both the mailbox database copy and the content index catalog copy. To seed just the mailbox database copy without seeding the content index catalog, use the *DatabaseOnly* parameter when running the [Update-MailboxDatabaseCopy](#) cmdlet. To seed just the content index catalog copy, use the *CatalogOnly* parameter when running the [Update-MailboxDatabaseCopy](#) cmdlet.

Seeding performance

Table 30 provides seeding performance information in our test environment. These results are based on the test environment configuration and can be different in the customer's environment. In our tests, the seeding was performed over 10 Gbps network with no latency between sites. Each source database was approximately 300 GB and the content index was about 35-40 percent. The total single replication source size (DB + Index) was approximately 420 GB.

Note:

In production environments, the content index is usually about 20 percent of the database size. It is 35-40 percent of the index size in our test environment due to the dynamic content generation feature in Loadgen.

Table 30. Seeding performance information

What is seeded?	Size	Average seeding time	Throughput (GB/min)
Single database (DB+index)	420 GB	2 hr 30 min	2.8 GB/min
Four databases (DB+index)	1,680 GB	18 hrs	1.5 GB/min

Environment validation with Loadgen

Overview

After completing the storage validation with Jetstress and determining that the storage is sized and performs as expected, the next step in the validation process is to use Loadgen tool to simulate MAPI workload against the entire Exchange infrastructure. Loadgen testing is necessary to determine how each Exchange component performs under a real, close-to-production user load.

Loadgen requires full deployment of the Exchange environment for validation testing. You should perform all Loadgen validation testing in an isolated lab environment where there is no connectivity to production data. Loadgen generates users and the workloads against entire Exchange environment including network and storage components.

Loadgen simulates the entire mail flow, helping to determine any bottlenecks in the solution. It is the only tool that helps you determine the CPU and memory resources that are necessary to sustain the load for which the Exchange environment was

designed.

In our tests, Exchange Server Load Generator 2010 (Loadgen) is used to simulate Outlook 2007 cache mode mailboxes with the following characteristics:

- The action profile is 100 messages per mailbox per day.
- Each mailbox is 600 MB in size.
- Each database contains 450 mailboxes.

To simulate a normal operation, the simulated workday duration is set to eight hours and each simulation runs for eight hours.

The 100-message profile sends 20 messages and receives 80 messages per mailbox per day. It is expected that during an eight hour simulated day the Mailbox server with 2700 active users will log approximately 1.88 sent messages per second and 7.5 delivered message per second. Use the formula below to calculate the expected number of sent and delivered messages per second.

$$\text{Messages sent per second} = \frac{(\text{number of mailbox users}) \times (20 \text{ messages sent})}{(8 \text{ hours} \times 60 \text{ minutes} \times 60 \text{ seconds})}$$

$$\text{Messages delivered second} = \frac{(\text{number of mailbox users}) \times (80 \text{ messages delivered})}{(8 \text{ hours} \times 60 \text{ minutes} \times 60 \text{ seconds})}$$

Peak load is used in this simulation for the 100 message Outlook 2007 Online Mode action profile. Peak load is enabled for an action profile by setting the simulated workday to four hours rather than eight hours. The 100-message action profile running in peak mode is expected to generate double the sent and delivered messages per second.

Use the following performance monitor counters on the Mailbox server to monitor the message sent and delivered rates:

MSExchangeIS Mailbox (_Total)\Messages Sent/sec

MSExchangeIS Mailbox (_Total)\Messages Delivered /sec

Mailbox server response times for client requests are tracked to determine the amount of time it takes the Mailbox server to respond to a client request. The response time average per request should not exceed 10 milliseconds. Use the following performance monitor counter on the Mailbox server to monitor response time.

MSExchangeIS\RPC Averaged Latency

As a best practice, disable Hyper Threading on the root server for all simulations for Exchange deployments by rebooting the server, entering the BIOS configuration, and disabling the Hyper Threading option.

The validity of the each test run is determined by comparing the results of select performance counters to a Microsoft specified criteria. Performance counter data is collected at 10-second intervals for the duration of each test run. The results of the first and last hours are discarded. Results are averaged over the remaining duration of test duration.

Table 31 lists the primary counters and their validation criteria:

Table 31. Primary counters and validation criteria

Performance monitor counter	Criteria
Processor(_Total)\% Processor Time	Not to exceed 80% during peak load
MSExchangeIS\RPC Averaged Latency	Not to exceed 10 msec.
MSExchangeIS Mailbox(_Total)\Messages Sent/sec	Approximately 0.002083 messages/second/mailbox
MSExchangeIS Mailbox(_Total)\Messages Delivered/sec	Approximately 0.008333 messages/second /mailbox
Logical Disk Disk Sec/read	Not to exceed 20 msec.
Logical Disk Disk Sec/write	Not to exceed 20 msec.

For additional information about monitoring Exchange 2010 performance and other key performance counters, visit "Performance and Scalability Counters and Thresholds" on Microsoft's Technet site at <http://technet.microsoft.com/en-us/library/dd335215.aspx>.

Table 32 lists the Loadgen tests that the team ran to measure the performance of the solution's Exchange Infrastructure.

Table 32. Environment validation tests with Loadgen

Test	Description
1	During normal operations at peak load, this test included a 100% concurrency test with 100 Outlook messages cached with a MAPI profile.
2	Within a site switchover and host failure, this test included 100% concurrency under peak load.
3	Site failure – During this test, Exchange VMs servicing DAG passive copies become active and provide service to users.

Test 1 – Normal operating condition – peak load

Test 1 objectives

In this test, the objective was to validate the entire Exchange environment under normal operating condition with the peak load. Each Hyper-V host and VMs performance were measured against Microsoft recommended performance targets and thresholds.

Test 1 configuration

In this test, all Exchange VMs were under normal operating condition. The Loadgen was configured to simulate peak load. The 100-message action profile running in peak mode was expected to generate double the sent and delivered messages generated per second.

Test 1 performance results and analysis

Table 33 and Table 34 show that all of the results achieved with Hyper-V servers and Exchange VMs were within the target metrics. During peak loads, the average CPU utilization on the primary Mailbox VMs was approximately 69 percent and the Hyper-V host utilization was about 70 percent. On the HUB/CAS VMs, CPU utilization was approximately 49 percent.

Table 33. Validation of expected load for Test 1

Parameter	Target	Tested 1 results
Message Delivery Rate / Mailbox	0.0056	0.0056
IOPs / Mailbox	0.12	0.22
Megacycles / Mailbox	1.35	2.25

Table 34. Performance results for Loadgen Test 1

	Performance counter	Target	Test 1 results
Hyper-V Root servers	Hyper-V Hypervisor Logical Processor(_total)\% Guest Run Time	<75%	66%
	Hyper-V Hypervisor Logical Processor(_total)\% Hypervisor Run Time	<5%	2%
	Hyper-V Hypervisor Logical Processor(_total)\% Total Run Time	<80%	68%
Primary Mailbox servers	Hyper-V Hypervisor Logical Processor(VP0-3)\% Guest Run Time	<80%	69%
	MSEExchange database\I/O database reads (attached) average latency	<20 ms	19 ms
	MSEExchange database\I/O database writes (attached) average latency	<20 ms <reads avg.	18 ms
	MSEExchange database\I/O log writes average latency	<20 ms	5 ms
	MSEExchange database\I/O log read average latency	<200 ms	5 ms
	MSEExchange\IS\RPC requests	<70	3
	MSEExchange\IS\RPC averaged latency	<10 ms	2 ms
Secondary Mailbox servers	Hyper-V Hypervisor Logical Processor(VP0-3)\% Guest Run Time	<80%	26%
	MSEExchange Database\I/O Database Reads (Recovery) Average Latency	<200 ms	0 ms
	MSEExchange Database\I/O Database Writes (Recovery) Average Latency	<200 ms	16 ms
	MSEExchange Database\I/O Log Read Average Latency	<200 ms	3 ms
	MSEExchange Replication(*)\ReplayQueueLength	<2	0
CAS/HUB servers combo	Hyper-V Hypervisor Logical Processor(VP0-3)\% Guest Run Time	<80%	53%
	MSEExchange RpcClientAccess\RPC Averaged Latency	<250 ms	8 ms
	MSEExchange RpcClientAccess\RPC Requests	<40	3
	MSEExchangeTransport Queues(_total)\Aggregate Delivery Queue Length (All Queues)	<3000	2.5
	MSEExchangeTransport Queues(_total)\Active Remote Delivery Queue Length	<250	0
	MSEExchangeTransport Queues(_total)\Active Mailbox Delivery Queue Length	<250	2.3

Test 2 – Within a site switchover and host failure – peak load

Test 2 objectives

In this test, the objective was to validate the entire Exchange environment under physical Hyper-V host failure/maintenance operating condition with the peak load. Each Hyper-V host and VMs performance were measured against Microsoft recommended performance targets and thresholds.

Test 2 configuration

During this test, two Hyper-V hosts within the site were shutdown. This resulted in a database switchover to a secondary local copy. During this switchover condition, one-half of the Exchange VMs handled double the load, which included 5,400 users. In this test, only half of the HUB/CAS servers processed Client Access and mail delivery.

Although the customer requirements included a design to sustain a single host failure, in this test we wanted to see the behavior of the extra load placed on HUB/CAS servers due to a double-host failure scenario. In this test, only two HUB/CAS servers serviced 5,400 active users.

Test 2 performance results and analysis

The results in Table 35 and Table 36 show that all of the Hyper-V servers and Exchange VMs that handled the peak load almost achieved the target metrics. There was slight overhead in CPU utilization (83 percent) on the primary Mailbox servers during the peak load because the generated load was two times what was expected.

Table 35. Validation of expected load for Test 2

Parameter	Target	Test 2 results
Message Delivery Rate / Mailbox	0.0056	0.0056
IOPs / Mailbox	0.12	0.24
Megacycles / Mailbox	2.4	2.7

Table 36. Performance results for Loadgen Test 2

	Performance counter	Target	Test 2 results
Hyper-V Root servers	Hyper-V Hypervisor Logical Processor(_total)\% Guest Run Time	<75%	66%
	Hyper-V Hypervisor Logical Processor(_total)\% Hypervisor Run Time	<5%	2%
	Hyper-V Hypervisor Logical Processor(_total)\% Total Run Time	<80%	68%
Primary Mailbox servers	Hyper-V Hypervisor Logical Processor(VP0-3)\% Guest Run Time	<80%	83%
	MSEExchange database\I/O database reads (attached) average latency	<20 ms	20.5 ms
	MSEExchange database\I/O database writes (attached) average latency	<20 ms <reads avg.	23 ms
	MSEExchange database\I/O log writes average latency	<20 ms	8 ms
	MSEExchange database\I/O log read average latency	<200 ms	9 ms
	MSEExchange\IS\RPC requests	<70	9
	MSEExchange\IS\RPC averaged latency	<10 ms	2 ms
Secondary Mailbox servers	Hyper-V Hypervisor Logical Processor(VP0-3)\% Guest Run Time	<80%	21%
	MSEExchange Database\I/O Database Reads (Recovery) Average Latency	<200 ms	0 ms
	MSEExchange Database\I/O Database Writes (Recovery) Average Latency	<200 ms	21 ms
	MSEExchange Database\I/O Log Read Average Latency	<200 ms	4 ms
	MSEExchange Replication(*)\ReplayQueueLength	<2	0
CAS/HUB server combo	Hyper-V Hypervisor Logical Processor(VP0-3)\% Guest Run Time	<80%	78%
	MSEExchange RpcClientAccess\RPC Averaged Latency	<250 ms	18
	MSEExchange RpcClientAccess\RPC Requests	<40	14
	MSEExchangeTransport Queues(_total)\Aggregate Delivery Queue Length (All Queues)	<3000	49
	MSEExchangeTransport Queues(_total)\Active Remote Delivery Queue Length	<250	0
	MSEExchangeTransport Queues(_total)\Active Mailbox Delivery Queue Length	<250	43

Test 3 – Site failure simulation

Test 3 objectives

The objective of this test was to validate the environment's performance when a third database copy is activated on passive Mailbox VMs. We wanted to measure the performance of the Hyper-V root server in the oversubscription scenario where the VP to LP ratio was 9:8. Each Hyper-V host and VMs performance were measured against Microsoft recommended performance targets and thresholds.

Test 3 configuration

During this test, a third database copy on the passive Mailbox VMs was activated. The client load was generated against all active copies simulating the load for 10,800 users. In this test, only half of the HUB/CUS servers were handling the client access and mail delivery for 10,800 users.

Test 3 performance results and analysis

The results in Table 37 and Table 38 show that all of the Hyper-V servers and Exchange VMs were able to handle the peak workload and the majority of the achieved results were within the target metrics. The solution team observed a slight increase in logical processor (LP) utilization of up to 85 percent on the Hyper-V root servers. This was due to the additional load placed on CUS and HUB servers as they were serving double the load of user mail activity.

Table 37. Validation of expected load for Test 3

Parameter	Target	Tested 3 results
Message Delivery Rate / Mailbox	0.0056	0.0056
IOPs / Mailbox	0.12	0.13
Megacycles / Mailbox	2.0	1.2

Table 38. Performance results for Loadgen Test 3

	Performance counter	Target	Test 3 results
Hyper-V Root servers	Hyper-V Hypervisor Logical Processor(_total)\% Guest Run Time	<75%	85%
	Hyper-V Hypervisor Logical Processor(_total)\% Hypervisor Run Time	<5%	2%
	Hyper-V Hypervisor Logical Processor(_total)\% Total Run Time	<80%	87%
Primary Mailbox servers	Hyper-V Hypervisor Logical Processor(VP0-3)\% Guest Run Time	<80%	63 %
	MSExchange database\I/O database reads (attached) average latency	<20 ms	16 ms
	MSExchange database\I/O database writes (attached) average latency	<20 ms <reads avg.	13 ms
	MSExchange database\I/O log writes average latency	< 20 ms	3 ms
	MSExchange database\I/O log read average latency	<200 ms	4 ms
	MSExchange\IS\RPC requests	<70	3
	MSExchange\IS\RPC averaged latency	<10 ms	2 ms
Secondary Mailbox servers (active in this test)	Hyper-V Hypervisor Logical Processor(VP0-3)\% Guest Run Time	<80%	21%
	MSExchange database\I/O database reads (attached) average latency	<20 ms	17 ms
	MSExchange database\I/O database writes (attached) average latency	<20 ms	12 ms
	MSExchange database\I/O log writes average latency	<20 ms	3 ms
	MSExchange\IS\RPC requests	<70	3
	MSExchange\IS\RPC averaged latency	<10ms	2 ms
CAS/HUB server combo	Hyper-V Hypervisor Logical Processor(VP0-3)\% Guest Run Time	<80%	63%
	MSExchange RpcClientAccess\RPC Averaged Latency	<250 ms	9 ms
	MSExchange RpcClientAccess\RPC Requests	<40	7
	MSExchangeTransport Queues(_total)\Aggregate Delivery Queue Length (All Queues)	<3000	5
	\MSExchangeTransport Queues(_total)\Active Remote Delivery Queue Length	<250	0
	\MSExchangeTransport Queues(_total)\Active Mailbox Delivery Queue Length	<250	4

Datacenter switchover validation

Datacenter switchover process

Managing a datacenter or site failure is different than managing the types of failures that can cause a server or database failover. In a high-availability configuration, the system initiates the automatic recovery, and the failure typically leaves the messaging system in a fully functional state.

By contrast, a datacenter failure is considered to be a disaster recovery (DR) event. When this happens, you must perform the recovery manually in order to restore the server and end the outage. This process is called a datacenter switchover.

As with many DR scenarios, prior planning and preparation for a datacenter switchover can simplify your recovery process and reduce the duration of your outage. For more information about datacenter switchovers, including detailed steps for performing a datacenter switchover, see Datacenter Switchovers at <http://technet.microsoft.com/en-us/library/dd351049.aspx>.

There are four basic steps that you need to complete to perform a datacenter switchover after making the initial decision to activate the second datacenter:

1. Terminate a partially running datacenter
2. Validate and confirm the prerequisites for the second datacenter
3. Activate the Mailbox servers
4. Activate the other server roles

The following sections describe these steps in more detail.

Step 1. Terminating a partially failed datacenter (DAG is in DAC mode)

When the DAG is in DAC mode, the specific actions to terminate any surviving DAG members in the primary datacenter are as follows:

The first step depends on the state of the failed datacenter. If the mailbox servers in the failed datacenter are still accessible (usually not the case), issue the following power shell cmdlet:

Stop-DatabaseAvailabilityGroup -ActiveDirectorySite <insertsitename>

If the Mailbox server is unavailable but Active Directory is operating in the primary datacenter, issue the following power shell cmdlet:

Stop-DatabaseAvailabilityGroup -ActiveDirectorySite <insertsitename> - ConfigurationOnly

Failure either to turn off the Mailbox servers in the failed datacenter or to successfully perform the **Stop-DatabaseAvailabilityGroup** command against the servers will create the potential for “split-brain syndrome” to occur across the two datacenters. You may need to individually turn off each computer through power management devices to satisfy this requirement.

Step 2. Validate and confirm the prerequisites for the second datacenter

The second datacenter must now be updated to represent the primary datacenter servers that have stopped. In the second datacenter run the following cmdlet:

```
Stop-DatabaseAvailabilityGroup -ActiveDirectorySite <insertsitename> - ConfigurationOnly
```

The purpose of this step is to inform the servers in the second datacenter about which Mailbox servers are available to use when restoring service.

Step 3. Activating Mailbox servers (DAG is in DAC mode)

Before activating the DAG members in the second datacenter, we recommend that you validate that the infrastructure services in the second datacenter are ready for messaging service activation.

When the DAG is in DAC mode, the steps to complete activation of the Mailbox servers in the second datacenter are as follows:

1. Stop the Cluster service on each DAG member in the second datacenter. You can use the **Stop-Service** cmdlet to stop the service (for example, Stop-Service ClusSvc), or use **net stop clussvc** from an elevated command prompt.
2. Activate the Mailbox servers in the standby datacenter by running the following power shell cmdlet:

```
Restore-DatabaseAvailabilityGroup -ActiveDirectorySite <insertsitename>
```

When this command completes successfully, the quorum criteria reduces to the servers in the standby datacenter. If the number of servers in that datacenter is an even number, the DAG switches to using the alternate witness server, as identified by the setting on the DAG object.

3. Activate the databases by running the following power shell cmdlet:

```
Get-MailboxDatabase <insertcriteriatoselectDBs> | Move-Active-mailboxDatabase - ActivateOnServer <DAGMemberinSecondSite>
```

Or, the following command

```
Move-Active-mailboxDatabase -Server <DAGMemberinPrimarySite> - ActivateOnServer <DAGMemberinSecondSite>
```

4. Check the event logs and review all error and warning messages to ensure that the secondary site is healthy. Follow up on and correct all issues prior to mounting the databases.
5. Mount the databases using the following power shell cmdlet:

```
Get-MailboxDatabase <DAGMemberInSecondSite> | Mount-Database
```

**Step 4.
Activating
Client Access
servers**

Clients connect to service endpoints to access the Microsoft Exchange services and data. Therefore, activating Internet-facing Client Access servers involves changing DNS records to point to the new IP addresses that need to be configured for the new service endpoints. Clients will then automatically connect to the new service endpoints in one of two ways:

- Clients can continue to try to connect, and should automatically connect after the TTL has expired for the original DNS entry, and after the entry is expired from the client's DNS cache. Users can also run the **ipconfig/flushdns** command from a command prompt to manually clear their DNS cache.
- Outlook Clients starting or restarting will perform a DNS lookup on startup and will get the new IP address for the service endpoint, which will be a Client Access server or array in the second datacenter.

To validate this scenario with Loadgen, perform the following steps:

1. Change the DNS entry for the Client Access Array to point to the VIP of the HWLB in the secondary site.
2. Run the **ipconfig /flushdns** command on all Loadgen servers.
3. Restart the **Loadgen Load**.
4. Verify that the CAS servers in the secondary site are now servicing the load.

Validating primary datacenter service restoration (failback)

Overview

Failback is the process of restoring service to a previously failed datacenter. The steps used to perform a datacenter failback are similar to the steps used to perform a datacenter switchover. A significant distinction is that datacenter failbacks are scheduled, and the duration of the outage is often much shorter.

It is important that you do not perform the failback until the infrastructure dependencies for Exchange are reactivated, functioning, stable, and validated. If these dependencies are not available or healthy, it is likely that the failback process will cause a longer outage, and it is possible that the process could fail altogether.

**Mailbox server
role failback**

The Mailbox server role (DAG is in DAC mode) should be the first role that has to failback to the primary datacenter. The following steps detail the Mailbox server role failback.

1. Reincorporate the DAG members into the primary site by running the following power shell cmdlet:

Start-DatabaseAvailabilityGroup -ActiveDirectorySite <insertsitename>

2. Verify the state of the database copies in the primary datacenter.

After the Mailbox servers in the primary datacenter are incorporated into the DAG, allow them to synchronize their database copies. Depending on the nature of the failure, the outage length, and the actions that the administrator

takes during the outage, this may require reseeding the database copies.

For example, if during the outage, you remove the database copies from the failed primary datacenter to allow log file truncation for the surviving active copies in the second datacenter, you will need to re-seed the database. Each database can individually proceed from this point forward. After a replicated database copy in the primary datacenter is healthy, it can proceed to the next step.

3. During the datacenter switchover process, the DAG was configured to use an alternate witness server. Reconfigure the DAG to use a witness server in the primary datacenter using the following power shell cmdlet.

Set-DatabaseAvailabilityGroup -Identity <DAGName> -WitnessServer <WitnessServer>

4. Dismount the databases that you are reactivating in the primary datacenter from the secondary datacenter using the following power shell cmdlet.

Get-MailboxDatabase <insertcriteria> | Dismount-Database

5. After the dismounting the databases, move the Client Access server URLs from the secondary datacenter to the primary datacenter by changing the DNS record for the URLs to point to the Client Access server or to the array in the primary datacenter.

Important:

Do not proceed to the next step until the Client Access server URLs have moved and the DNS TTL and cache entries are expired. Activating the databases in the primary datacenter prior to moving the Client Access server URLs to the primary datacenter results in an invalid configuration (for example, a mounted database that has no Client Access servers in its Active Directory site).

6. You can now activate the databases by running the following power shell cmdlet:

Get-MailboxDatabase <insertcriteriatoselectDBs> | Move-Active-mailboxDatabase - ActivateOnServer <DAGMemberinSecondSite>

Or

Move-Active-mailboxDatabase -Server <DAGMemberinPrimarySite> - ActivateOnServer <DAGMemberinSecondSite>

7. Mount the Databases using the following power shell cmdlet:

Get-MailboxDatabase <insertcriteriatoselectDBs> | Mount-Database

Chapter 8: Conclusion

More and more IT organizations are realizing the benefits of virtualization. They are reducing costs and providing more agile IT services to their business. But with business-critical applications like Microsoft Exchange, these transitions take a great deal of time. In addition, users are often reluctant to try new storage strategies. In these situations, organizations not only want to know that it *will* work but also *how* it will work. Before making your next investment, you need a proven architecture with tested best practices.

The synergies created by the Microsoft, Cisco, and EMC partnership provide the assurance you need to move forward with virtualization in your environment. You can now reduce the risks and accelerate deployment with this jointly tested architecture.

The testing performed by Microsoft, Cisco, and EMC lowers customer's risk. By leveraging these best practices and lessons learned customers can reduce the chance that components will not work or that they have the wrong driver levels, or that the network or storage was not setup correctly.

The solution team rebuilt a sample customer environment to test these solutions. We also pushed the environment to its limits by stressing out the configuration with I/O workload generators and ran a series of failover scenarios to see how the environment would fail over. Since the results were so astounding, Microsoft, Cisco, and EMC published their findings to prove that their solutions can help businesses take the next steps in changing their storage environments.

	Before virtualization	After virtualization
Servers	Deployed in a physical fashion, customers are stuck in a cycle of: 1. Increasing workloads/data 2. Adding servers to satisfy growth 3. Adding people to manage complex solutions	Cisco UCS, deployed with Windows Server 2008 R2 with Hyper-V, enables faster provisioning and centralized management, which leads to increased efficiency and lower operational costs.
Network	Each server has its own independent network—setup, configuration, and maintenance can become challenging	The Cisco UCS architecture balances simplicity, performance, and density allowing you to centralize management and gain access to a shared 10GbE backplane, which provides better management and higher performance for reseeding, failovers, and failbacks.
Storage	Over-provisioned direct-attached configurations that must be maintained separately. Storage allocation, backup, replication, and archiving are all configured independently.	EMC Virtual Provisioning is built into the EMC Unified CX4-480. This allows customers to start small and expand into larger mailbox sizes as they grow. Storage area networks allow you to pool your storage, EMC virtual provisioning technology enables you to start with a small storage environment and pay as you grow. You can also centralize all infrastructure operations (storage, backup, replication, and archiving), allowing you to become much more efficient.

Appendixes

Appendix A: References

White papers

For additional information, see the white papers listed below.

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

Product documentation

For additional information, see the product documents listed below:

- **EMC CLARiiON CX4-480**
<http://www.emc.com/products/detail/hardware/clarion-cx4-model-480.htm>
 - **Cisco UCS Systems**
<http://www.cisco.com/en/US/netsol/ns944/>
 - **Microsoft Exchange Server 2010**
<http://technet.microsoft.com/en-us/library/bb124558.aspx>
-

Other documentation

For additional information, see the documents listed below.

- **Microsoft Exchange 2010 Mailbox Server Processor Capacity Planning**
<http://technet.microsoft.com/en-us/library/ee712771.aspx>
 - **Microsoft Exchange 2010 – Understanding Mailbox Database Cache**
<http://technet.microsoft.com/en-us/library/ee832793.aspx>
 - **Microsoft Exchange 2010 Mailbox Server Role Requirements Calculator**
<http://msexchangeteam.com/archive/2009/11/09/453117.aspx>
 - **Exchange 2010 Performance and Scalability Counters and Thresholds**
<http://technet.microsoft.com/en-us/library/dd335215.aspx>
 - **Microsoft Exchange 2010 Overview of Mailbox Server Role**
<http://technet.microsoft.com/en-us/library/bb124699.aspx>
 - **SPEC CPU2006**
<http://www.spec.org/cpu2006/>
 - **Intel Xeon Processor 5500 Series**
http://www.intel.com/Assets/en_US/PDF/prodbrief/xeon-5500.pdf
-

Appendix B: Terminology

The following table defines the terms used in this document.

Term	Definition
Bandwidth	The amount of data a storage system can process over time, which is measured in megabytes per second.
Converged Network Adapter	An adapter that presents network interface card (NIC) interfaces and host bus adapter (HBA) interfaces to an operating system. This adapter provides server connectivity for data, storage and management resources using an Ethernet unified fabric.
Disk Transfers/sec	The rate of read and write operations on the disk.
EMC Unisphere	Software that provides the next generation of storage management and presents a single, integrated, and simple web-based interface for the CLARiiON and Celerra storage systems.
Exchange user profile types	<p>The Exchange user profile includes the message profile, the mailbox size, and the number of users to be included in the Exchange environment.</p> <p>For more information about Exchange user profile types, visit Microsoft at http://technet.microsoft.com/en-us/library/ee712771.aspx.</p>
Fibre Channel over Ethernet (FCoE)	Encapsulation of Fibre Channel frames over Ethernet networks. This allows Fibre Channel to use 10 Gigabit Ethernet networks while preserving the Fibre Channel protocol.
FLARE LUN	A logical unit number (LUN) that is created from a CLARiiON RAID group and is not thin provisioned.
Jetstress	Microsoft Exchange Server Jetstress 2010 is a tool used to verify the performance and stability of a disk subsystem prior to putting a Microsoft Exchange Server into production. Jetstress helps to verify disk performance by simulating the Exchange disk input/output (I/O) load. Specifically, Jetstress simulates the Exchange database and log file loads produced by a specific number of users.
Logical Processor (LP)	The representation of a physical processor core of a multi-core processor.
Logical unit number (LUN)	A unique identifier used to identify logical storage objects in a storage system.
Mailbox server building block	EMC best practices for designing an Exchange 2010 Mailbox server and storage. A building block represents the required resources required to support a specific number of Exchange 2010 users on a single Exchange Mailbox server VM.
Pool LUNs	<p>A pool LUN consists of two types, <i>thin LUN</i> and newly-introduced <i>thick LUN</i> (fully provisioned LUN):</p> <ul style="list-style-type: none">• A fully provisioned thick LUN is a logical unit of storage created within a storage pool and consumes physical

	<p>storage of the LUN size up front from the storage pool. You can shrink or expand the fully provisioned LUN.</p> <ul style="list-style-type: none"> • A thin LUN consumes physical storage from the storage pool only when data is written to it. This thin LUN competes with other thin LUNs in the pool for available storage.
Processor (CPU) Core	A processor is the unit that reads and executes program instructions. Processors were originally developed with only one core. The core is the part of the processor that actually performs the reading and executing of the instruction. A multi-core processor is a processing system composed of two or more independent cores.
RAID 10	A RAID method that provides data integrity by mirroring data onto another disk. This RAID type provides the greatest assurance of data integrity at the greatest cost in disk space.
RAID 5	A RAID method where data is striped across disks in large stripes. Parity information is stored so data can be reconstructed, if necessary. In this situation, one disk can fail without incurring any data loss. Performance is good for read activity, but slower for write activity.
Response time	The interval of time between submitting an I/O request and receiving a response.
Service Profile	The service profile maintains configuration information about the server hardware, interfaces, fabric connectivity, and server and network identity. This information is stored in a format that you can manage through Cisco UCS Manager. All service profiles are centrally managed and stored in a database on the fabric interconnect.
Short-stroked drives	A technique where data is laid out on partially populated disks to reduce the spindle head movement and to provide higher IOPS at a very low latency.
Soft Zoning	Zoning implementation that is based on fabric device names such as WWPN and WWNN.
Storage pools	<p>A set of disks with the same redundancy (RAID 5, RAID 6 or RAID 1/0 only), and shares its user capacity with one or more pool LUNs.</p> <p>EMC storage pools comprise of:</p> <ul style="list-style-type: none"> • Homogeneous pools that contain a single drive type, that is, only EFD, FC, or SATA. Auto-tiering is not available. • Heterogeneous pools that contain mixed drive types of EFD, FC, and SATA. Only two different drive types are needed for FAST to enable auto-tiering.
Storage processor (SP)	On a CLARiiON storage system, it is a circuit board with memory modules and control logic that manages the storage-system I/O between the host's FC adapters and the disk modules.
Throughput	The number of individual I/Os the storage system can process over time, which is measured in I/Os per second.
Unified Fabric	Consolidation of data, storage, and management traffic onto a single, general-purpose, high-performance, highly available network to greatly simplify infrastructure and reduce costs.

Virtual Processor (VP or vCPU)	Virtual processor defines the number of processors that are presented to the operating system that is running in a virtual machine.
Write penalty	<p>The write penalty is inherent in RAID data protection techniques, which require multiple disk I/O requests for each application write request.</p> <p>For RAID 10, two I/O requests are required for each application request; for RAID 5, four I/O requests are required for each application request.</p>
Zoning	Partitioning of a Fibre Channel fabric into smaller subsets to restrict interference, add security, and to simplify management.