

Microsoft® Exchange Server 2010 Implementation on Dell™ Active System 800v

*A Design and Implementation Guide for Exchange Server 2010 on
Active System 800 with VMware vSphere*

Dell Global Solutions Engineering

Revision: A01

February 2013



This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

© 2012 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. *Dell*, the *Dell logo*, *PowerEdge*, *Force10*, *EqualLogic*, and *OpenManage* are trademarks of Dell Inc. *Intel* and *Xeon* are registered trademarks of Intel Corporation in the U.S. and other countries. *Microsoft*, *Windows*, and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *VMware*, *vCenter*, *vSphere*, *ESXi*, *vCloud*, and *vMotion* are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

Rev A01

Contents

Executive Summary	6
Audience and Scope	7
Introduction	7
Overview of Active System 800.....	7
Overview of Microsoft® Exchange Server 2010	9
Dell Services Offerings	10
Designing Exchange 2010 on Active System 800.....	10
Active System 800 Features for Exchange Server 2010	11
Designing the Storage	12
Guidelines for Shared Storage in Virtualized Environments	12
Guidelines for Exchange Server Storage	13
Designing Exchange Specific Networks.....	13
Sizing Virtual CPU Requirements.....	14
Virtual CPU Requirements for the Mailbox Server Role	14
Virtual CPU Requirements for the Hub Transport and Client Access Server Roles	15
Sizing the Memory Requirements	15
Memory Requirements for Mailbox Server	16
Memory Requirements for the Hub Transport and Client Access Server Roles	16
Reference Implementation for 5,000 Mailboxes on Active System 800v.....	17
Sizing and Configuring Storage	18
Capacity Requirements per Copy.....	18
IOPS Requirement per Copy	18
DAG Storage Configuration	19
Identifying and Configuring Network Requirements	21
Host Network Configurations.....	22
VM Network Configurations.....	24
Identifying and Configuring CPU Requirements	24
Identifying and Configuring Memory Requirements.....	25
Configuring VM Attributes	25
VM Placements	28
Solution Summary	30
Reference Implementation Verification of 5,000 Mailbox Users	31
Performance Verification.....	31
Verifying Storage Requirements Using Jetstress.....	32

Verifying the End-to-End Performance of Exchange Server with Active System 800v	32
Verifying the VM Placements	37
Assessing the Impact of Single DAG Node Failure	38
Failure Scenario	38
Database Redistribution	39
Summary	40
References	41
Appendix A: Sizing Virtual CPU for Mailbox Server Role	42
Appendix B: Sizing Virtual CPU for Hub Transport and Client Access Server Roles	44
Appendix C: Sizing Virtual CPU for Other Exchange Server Roles.....	45
Appendix D: Sizing Memory for Mailbox Server Role	46
Summary of Memory Requirements for Various Exchange Server Roles.....	47

Tables

Table 1. Active System 800v Configurations	9
Table 2. Considerations for Number of Virtual Machines per LUN.....	13
Table 3. Recommended Memory Configurations for Hub Transport and Client Access	16
Table 4. Example Solution Requirements for Exchange Server	17
Table 5. Resource Requirements for Exchange Server Solution	26
Table 6. Solution Summary	30
Table 7. Mailbox Profile for Jetstress Test	32
Table 8. Performance Test Plan for Exchange Load Generator.....	33
Table 9. Test Results for VMs Running Mailbox Server Role in a Steady State Configuration.....	34
Table 10. Test Results for VMs Running HUBCAS Server Roles in a Steady State Configuration	35
Table 11. Test Results for VMs Running Mailbox Server Role in a Failed Host Configuration	36
Table 12. Test Results for VMs Running HUBCAS Server Roles In a Failed Host Configuration.....	37
Table 13. Steady State and Failed DAG Configurations.....	38
Table 14. Per Mailbox Database Cache, IOPS and CPU Estimates.....	42
Table 15. Mailbox To (Hub Transport + Client Access) Processor Core Ratio	44
Table 16. Processor Configurations for Various Exchange Server Roles.....	45
Table 17. Database Cache Requirements.....	46
Table 18. Mailbox Memory Recommendations by Microsoft Based on Database Cache	46
Table 19. Memory Recommendations	47

Figures

Figure 1.	High-level Architecture of Active System 800	8
Figure 2.	A Logical Layout for the DAG Storage	19
Figure 3.	DAG State After a Storage Array Failure	20
Figure 4.	DAG State After Server Failure	21
Figure 5.	Active System 800v NPAR Configuration for Exchange Server.....	22
Figure 6.	Sample Implementation of Host Network Configuration	23
Figure 7.	VM Logical Networking	24
Figure 8.	Allocating RDM Disks to Virtual SCSI Controllers.....	27
Figure 9.	VM Attributes for Mailbox Server Role.....	28
Figure 10.	VM Placements.....	29
Figure 11.	Test Bed for Exchange Server Verification with Active System 800v	31

Executive Summary

Traditional Information Technology (IT) infrastructures have been transformed with the use of virtualization technologies in the datacenter. Some of the benefits of virtualization, including consolidation and availability, have enabled creation of large clusters of hardware resources managed with virtualization and related software. These architectures improve the ease of management of underlying hardware resources by administrators and enable coexistence of different tiers of enterprise applications on the same hardware.

Dell™ Active Infrastructure is a family of converged infrastructure solutions that combine servers, storage, networking, and infrastructure management into an integrated and optimized system that provides virtualized resource pools. Active Infrastructure helps IT rapidly respond to dynamic business demands, maximize data center efficiency, and strengthen IT service quality. Dell Active System 800 is a pre-integrated virtualized infrastructure system that enables IT and business stakeholders to accelerate application and IT service delivery. Designed for reliability, Active System 800 eliminates single points of failure and incorporates redundancy into all critical infrastructure components.

Critical enterprise applications including Microsoft Exchange Server 2010 were designed to be run on dedicated physical hardware with best practices around performance and availability detailed in Microsoft and partner documentation. With the introduction of virtualization in the datacenter some best practices for moving these applications into virtualized servers running multiple virtual machines (VMs) were also provided. However many approaches focused on dedicating a set of hardware resources for the application. This ensured the appropriate performance and availability of the application during hardware failure scenarios. However it also led to the creation of silos of hardware resources dedicated for specific application. This has led to increase in management complexity for IT administrators.

This guide presents a framework for implementing Microsoft Exchange Server 2010 solution on Active System 800. This framework builds on the architecture best practices of the Active System and uses VMware virtualization software capabilities to develop rules that allow application VMs to be deployed and managed without administrators worrying about the exact physical servers they might be running on. The design also provides dedicated storage resources for application data, guaranteeing adequate storage performance to meet application requirements.

The Design and Implementation Guide also provides an instantiation of the application design to support up to 5,000 users using Active System 800. This implementation was then verified for performance and availability by running the Exchange Loadgen and Jetstress tools to determine steady-state and degraded-state performance characteristics and ensuring that the virtualization resource scheduling features implemented for the application VMs lead to the application returning to steady-state quickly upon hardware resource failure.

Audience and Scope

This guide is intended for IT professionals and administrators interested in designing and deploying custom sized Exchange 2010 solution up to 5,000 mailboxes, on Dell Active System 800 with VMware vSphere. While the guide provides an overview of Active System 800 and Exchange Server, the reader is expected to have certain understanding of VMware vSphere 5.1 and Microsoft Exchange Server 2010 environment, as well as Active System 800.

Introduction

Recent advancements in server, storage, and networking technology have enabled building efficient infrastructure solutions for the private cloud. Dell's latest hardware portfolio is well equipped to take advantage of the technology advancements and help build solutions with converged network infrastructure. Active System 800 is one of such systems that include the latest Dell PowerEdge™ blade servers, 10 GbE switching, and Data Center Bridging (DCB) and Dell EqualLogic™ iSCSI storage to build private cloud solutions.

Virtualization software from VMware®, including vSphere™ 5 with vCenter™ Server, provides fine-grained virtual resource management tools that enable advanced management of application resources to provide reliable performance guarantees and increased availability without tying virtual machines (VMs) to individual servers. While the underlying infrastructure is becoming more feature-rich, accurately sizing and verifying Exchange Server 2010 on such an infrastructure requires a new approach. Some of the considerations include accurately defining resource requirements and maintaining application best practices at the same time, while guaranteeing application performance within the thresholds that are prescribed by Microsoft. As a consequence, a methodology that correctly designs and deploys Exchange Server 2010 on new converged infrastructure is essential.

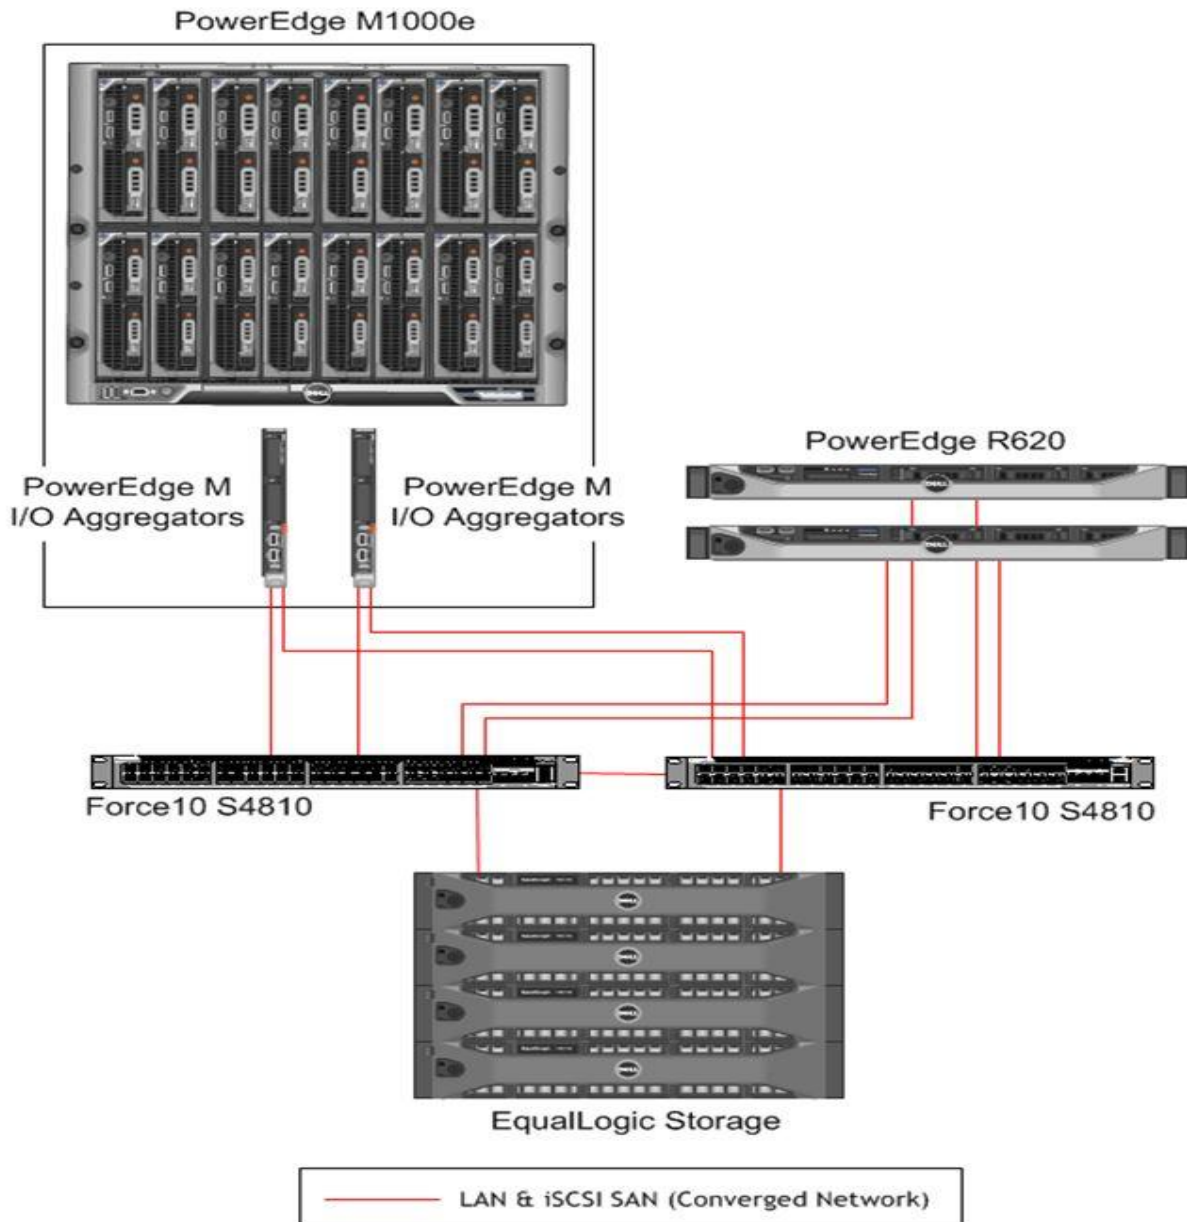
One such methodology is described in this guide. The methodology helps reduce dependencies on the underlying infrastructure from a sizing and deployment perspective, while simplifying the management of Exchange Server 2010 ecosystem. It does this by abstracting out the Exchange solution in the form of a virtualized resource configuration consisting of memory, virtual Central Processing Units (vCPUs), and virtual network adapters that are provided by Active System 800. The resource configuration also consists of an appropriate number of VMs pertaining to various server roles that are part of the Exchange server deployment. The performance of Exchange Server 2010 is guaranteed as long as the underlying platform is able to meet the virtualized resource configuration requirements and resource allocation policies.

Overview of Active System 800

Converged Infrastructure is the latest advance in private and public cloud infrastructure architectures. Converged infrastructure based platforms aggregate various information technology components into a single, optimized computing platform. Typically, the components of a converged infrastructure based platform includes servers, data storage devices, networking equipment and software for infrastructure management, monitoring and automation, thus creating a feature-rich infrastructure for the enterprise workloads. The Active System 800 is one such enterprise infrastructure system that has been designed and verified by Dell™ Engineering. It is available to be fully racked, cabled, and delivered to the customer's site to speed up the deployment.

Active System 800v includes Dell PowerEdge™ M1000e blade chassis with Dell PowerEdge™ M I/O Aggregator, Dell PowerEdge™ M620 blades, Dell EqualLogic™ Storage, Dell Force10™ network switches, and VMware vSphere 5.1. The system also includes Dell PowerEdge™ R620 servers as management servers. Management related software such as VMware vCenter Server, EqualLogic Virtual Storage Manager for VMware, and OpenManage™ Essentials, are included with the system. A high-level architectural layout of Active System 800v is shown in Figure 1.

Figure 1. High-level Architecture of Active System 800



Active System 800 is offered in configurations with either VMware® vSphere® (Active System 800v) or Microsoft® Windows Server® 2012 with Hyper-V® role enabled (Active System 800m) hypervisors.

As a converged infrastructure system, Active System 800 employs the Data Center Bridging (DCB) suite of standards and technology available on the Dell Force10 network infrastructure to prioritize different traffic types in a converged fabric. Active System 800v converged infrastructure is designed to enable Ethernet transport reliability for application workloads. The underlying design principles make Active System 800v a suitable converged infrastructure system for virtualizing I/O intensive workloads. The table below highlights the different offerings for Active System 800v by Dell.

Table 1. Active System 800v Configurations

	8 Blades Configuration	16 Blades Configuration	24 Blades Configuration	32 Blades Configuration
Compute capacity (Total number of cores)	128	256	384	512
Memory capacity (in TB)	1	2	3	4
Storage capacity (Raw TB)	43.2	86.4	129.6	172.8
Number of drives	(2x24) 10K 900GB 2.5" SAS drives	(4x24) 10K 900GB 2.5" SAS drives	(6x24) 10K 900GB 2.5" SAS drives	(8x24) 10K 900GB 2.5" SAS drives

As can be seen in Table 1, Active Systems present increasingly larger amounts of compute, memory, and storage capacity in moving from the configurations consisting of eight blades to the thirty two blades. For the purpose of the reference architecture described in this guide, Active System 800v is considered as the resource repository. This repository allows IT Administrators to run Exchange Server, Lync Server, and SharePoint Server along with other workloads, beyond just the sample configuration presented in this guide.

For more information on Active System 800v specifications, please refer the Active System 800v documentation.

Overview of Microsoft® Exchange Server 2010

Exchange Server 2010 is the leading enterprise messaging system, comprised of multiple sub-systems also known as server roles. A server role is an application layer entity, and multiple roles can be collocated on a single machine to provide certain services. Here is a quick overview of the Exchange 2010 server roles:

- **Mailbox Server Role (MBX):** A back-end server capable of hosting mailboxes and public folders. Multiple MBX roles can be clustered using a Database Availability Group (DAG).
- **Client Access Sever (CAS):** A server role that supports all messaging clients, such as Outlook, mobile cellular devices and Exchange Web Services.
- **Hub Transport Server (Hub):** A routing server that routes a message within the Exchange organization.

- **Edge-Transport Server (Edge):** A server role residing on the edge of the topology that routes messages in and out of the Exchange organization.
- **Unified Messaging Server (UM):** A server role that connects a PBX system to Exchange topology and helps combine voice and email messages into a single messaging infrastructure.

NOTE: This guide concentrates primarily on the Mailbox, Client Access, and Hub Transport server roles in the contexts of the design framework and the reference implementation. Exchange clients corresponding to mobile cellular devices are not considered in the study.

In order to provide continuous scalability and availability, Exchange Server 2010 uses Database Availability Groups. A DAG is a group of two or more (up to 16) servers hosting mailbox databases to provide database level protection and recovery. The native high availability leverages the cluster technologies in Continuous Cluster Replication (CCR) based on a Majority Node Set (MNS) model of Microsoft Cluster Service (MSCS). There have been significant improvements with Exchange Server 2010. For example, it has been optimized to take advantage of low-cost storage. With the enhanced Exchange Store schema, the IOPS requirements have been significantly reduced compared to the previous versions. The database pages in data structures (B-Tree) are stored in a more contiguous fashion, as opposed to random allocation as in previous versions of Exchange. This contiguity provides a sequential pattern for reads and writes and allows more efficient I/O to and from a disk subsystem. The database page size has also been increased from 8KB (in Exchange 2007) to 32KB. Multiple database page operations are coalesced to produce a single large I/O operation. All these factors, including larger I/O and page size, have contributed towards reducing the overall IOPS in Exchange Server 2010.

Dell Services Offerings

Dell Services provides customers a portfolio of professional services including deployment, consulting and support. Deployment services for Active System 800 helps customers deploy and configure the solution ready to be deployed into their datacenter. Dell Consulting for Exchange provides expert guidance and proven processes to help reduce the time, costs, and risk of transitioning to latest Exchange solution. Dell Consulting has invested in defining repeatable consulting services, which will ensure email and calendaring platforms based on Exchange are optimized, flexible, efficient and secure. Dell ProSupport™ for Exchange provides customers with a consolidated, single point of contact for their software support needs and ensures solution availability. ProSupport allows customers unlimited access to Dell's 24x7 chat, email and phone support services to help with configuration issues, how-to-assistance, critical patches, updates and disaster recovery.

For detailed services information, please visit dell.com and contact Dell Services representatives.

Designing Exchange 2010 on Active System 800

This section describes a stepwise approach along with generic guidelines to design solutions for Exchange Server 2010 deployments on a completely virtualized and converged infrastructure such as Active System 800v. The goal of the methodology and guidelines is to design an Exchange Server 2010 solution that is independent of the underlying virtual platform. Each Exchange Server 2010 server role requires distinct system resources and needs to be sized with attention to role specific demands. Because we are considering a completely virtualized computing platform based on VMware vSphere 5.1

to run the Exchange Server workload, this section considers virtualization best practices from VMware and application best practices from Microsoft. These guidelines primarily consider virtual processors, memory and storage sizing, and configurations with respect to various server roles included in a typical Exchange Server 2010 deployment.

Active System 800 Features for Exchange Server 2010

Before diving into the design framework, it would be worthwhile summarizing the services and the features Active System 800v has to offer for Exchange Server 2010. Exchange Server 2010 is an I/O intensive, enterprise-level messaging workload that needs resource guarantees along with ease of resource management.

Benefits of Active System 800 include:

- **Power efficiency:** PowerEdge M1000e chassis and PowerEdge M620 blade servers require less power and cooling than alternatives, resulting in reduced data center power and cooling costs.
- **Converged fabric:** PowerEdge M I/O Aggregator enables interoperability with a non-Dell network and allows the server administrator to push blade chassis networking features to the top-of-rack switch. Additionally, Data Center Bridging (DCB) over iSCSI enables single converged network fabric for storage and LAN traffic. It supports intra-chassis networking (east/west system communication) with 10GbE throughput.
- **Consistent hardware, firmware, and software configurations** across cluster hosts are used to avoid performance variations for similar VMs running on different VMware ESXi hosts.
- **Rapid and simple scalability:** Capacity can be added in a modular, non-disruptive manner for improved high availability for Exchange DAG.

Deploying Exchange server in Active System 800 environment can be very beneficial. However, it requires that Exchange server resource requirements be sized based on virtual resource requirements, which demands a new design and implementation methodology. A structured approach is discussed in the following sub-sections.

Designing the Storage

This section describes how to configure shared storage for Exchange server roles, guest Operating System (OS) images, and databases. Because Exchange server is a storage-intensive enterprise workload, identifying the storage requirements for Exchange databases is one of the critical steps in this design framework. Storage sizing for Exchange primarily consists of two aspects:

- Capacity sizing: Determines how many mailboxes of a given size can be incorporated in a given storage space.
- Sizing for IOPS requirements: Determines the type of mailbox profile that can be supported on the storage.

For details on storage sizing for Exchange databases, refer to the *'Mailbox role storage capacity calculations'* section in the Introduction to Microsoft Exchange Server 2010 sizing¹ paper. The following subsections provide guidelines on configuring storage in a virtual environment and summarize guidelines for sizing Exchange database storage.

Guidelines for Shared Storage in Virtualized Environments

- Establishing storage multi-paths is essential to improve the performance of Exchange databases. VMware recommends setting up at least four paths from the ESXi host to the storage array, which requires at least two HBA adapters to be configured per ESXi host.
- Raw Device Mapping (RDM) is recommended for Exchange databases because RDM enables the database portability feature of Exchange databases. RDM volumes are supported for both virtual machines and physical servers. It should be noted that RDM also works with third-party clustering.
- The VMFS is recommended for guest OS store (VM store). VMware template cloning is one of the benefits of using VMFS for Exchange VMs. A mixed configuration is recommended, where the guest OS is installed on VMFS volumes and the Exchange database and log files use RDM.
- There are tradeoffs in mapping one VM to one LUN, as compared to having many VMs mapped to one large LUN. Table 2 discusses some of those tradeoffs.
- VMware provides up to four virtual SCSI adapters and up to 16 disks per virtual SCSI adapter per VM. Balanced use of separate virtual SCSI adapters for mapping Exchange databases is highly recommended. This means that each virtual SCSI adapter of the VM should be mapped to an equal number of active and passive storage volumes.
- Snapshots should be disabled and Data Center Bridging should be enabled per EqualLogic storage member.

¹Introduction to Exchange Server Sizing:

<http://i.dell.com/sites/content/business/solutions/whitepapers/en/Documents/exchange-server-2010-sizing.pdf>

Table 2. Considerations for Number of Virtual Machines per LUN

Single LUN (VMFS/RDM) mapped to single VM	One VMFS volume sharing multiple VMs
Utilization is not optimal	Utilization of the volume is better
More management efforts are required	Provisioning and management of the volume as well as VM is easier
Results in islands of allocations	The I/O performance may get impacted adversely
Good for I/O heavy workloads	

NOTE: Active System 800v ensures to facilitate the storage best practices for Exchange Server 2010.

Guidelines for Exchange Server Storage

Configuring storage pertaining to Exchange database is critical to the overall performance of the application. Here are the guidelines for configuring storage for an Exchange solution:

- Isolate the Microsoft Exchange Server database and log workload from other I/O intensive applications; assign the Exchange workload to its own set of disks.
- Deploy each DAG copy on its own set of physical disks and separate storage controller; place active and passive copies of the database in separate volumes.
- Design the storage for the active and passive copies of the mailboxes to be identical in terms of capacity and performance.
- Virtual disks that dynamically expand are not supported by Exchange.
- NTFS allocation unit size should be set to 64KB for optimal performance of the volume.
- Physical disk cache should be disabled. This is ensured by Active System 800v.
- Microsoft does not support using NFS store for Exchange databases.

NOTE: The administrator must follow these guidelines while configuring and deploying the Exchange solution on Active System 800v.

Designing Exchange Specific Networks

Once the storage has been sized and configured, the next step is to identify the specific network requirements. Because Exchange is an enterprise-level messaging application, it requires very specific allocation of network resources and bandwidth guarantees. This section provides design guidelines for network resource allocations for VMs running various Exchange server roles:

- The DAG replication traffic must be separated from the client Messaging Application Programming Interface (MAPI) traffic at the DAG node, on different physical/virtual network adapters; it can share the same physical/virtual switch with different VLANs.
- For Exchange databases residing on SAN, it is recommended to configure separate virtual switches for the LAN traffic and the SAN traffic (iSCSI), at the hypervisor level.

- Implement redundant components (switches, network adapters) to provision a resilient network infrastructure between the endpoints (stack, LAG, load balancing, or network card teaming).
- It is recommended that DCB is enabled for the switch ports hosting the EqualLogic PS Series array controller connections, and the network adapters on the hosts are dedicated to converged network traffic.
- Ensure that jumbo frames (large Maximum Transmission Unit (MTU)) are enabled for the virtual switch ports interfacing the PS Series array controller connections.

NOTE: Active System 800v complies with the networking guidelines described above.

Considering the above recommendations, the design methodology prescribes having separate virtual network adapters for Exchange MAPI and DAG replication networks.

Sizing Virtual CPU Requirements

The next step is to address Exchange server role-specific computing demands. VMs with more than one vCPUs are considered as symmetric multi-processing (SMP) VMs. VMs with SMP exhibit enhanced performance, because a single VM can use multiple physical processor cores simultaneously. VMs running on a vSphere 5.1 host can be configured with up to 64 vCPUs, which indicates the feasibility of running compute intensive enterprise workloads. The biggest advantage of a SMP VM is its ability to use multiple processors to execute multiple tasks concurrently, thereby increasing throughput, especially for workloads like Exchange, that use parallelization. The following best practices must be taken in account, while considering VMs with multiple virtual processors:

- Allocate multiple vCPUs to a VM if the anticipated Exchange workload can truly take advantage of all the vCPUs.
- If the exact workload is not known, size the VM with a smaller number of vCPUs initially and increase the number later if necessary.
- For performance critical Exchange VMs (i.e., production systems), try to ensure the total number of vCPUs assigned to all the VMs is equal to or less than the total number of cores on the ESXi host machine.

In order to take advantage of Symmetric multi-threading (SMT), it is recommended to have Hyper-threading enabled on the ESXi hosts running Exchange specific VMs. Essentially, Hyper-threading technology enables a single physical core to act like two logical cores, which in turn improves the processor performance. An ESXi host running on an 8-core server with Hyper-threading enabled will perceive 16 total logical cores. It is recommended that in a virtualized environment, the ratio of vCPUs to logical cores should be 1:1 when CPU hyper-threading is enabled. When hyper-threading is disabled, the ratio of vCPUs to physical cores should also be 1:1. In order to guarantee CPU performance, Microsoft prohibits over-subscription of processing capabilities. A 10% overhead on the CPU should be taken into account for virtualized environments. It should also be noted that the enhancements provided by virtualized execution environments, such as dynamic memory reservation and CPU affinity, should be disabled while running Exchange mailbox server.

Virtual CPU Requirements for the Mailbox Server Role

The Mailbox server role deals directly with the storage database and require more computing power as compared to other Exchange server roles. Megacycles are used to measure the processor capacity. Each

Exchange mailbox user variably consumes the processing capacity of the server, depending upon the mailbox profile and activity. Standard steps can be followed to identify the CPU requirements for the mailbox server role. Please refer to ‘*Mailbox role megacycles calculations section*’ in the Introduction to Microsoft Exchange Server 2010 sizing² paper and Appendix A: Sizing Virtual CPU , for more details.

Virtual CPU Requirements for the Hub Transport and Client Access Server Roles

The recommended processor core ratio of Mailbox to (Hub Transport + Client Access) server roles is 1:1. For details on sizing the CPU for various configurations of Hub Transport and Client Access server roles, please refer to Appendix B: Sizing Virtual CPU for Hub Transport and Client Access Server Roles. To size the CPU for server roles other than Mailbox, Hub Transport and Client Access, please refer Appendix C: Sizing Virtual CPU for Other Exchange Server Roles.

To implement the recommended sizing for the virtual processors, an ability/service to guarantee exact CPU resources is needed from the underlying platform. As a result, it is recommended to make use of VM configuration parameters such as Reservations, Shares, and Limits to guarantee the processing resources.

Sizing the Memory Requirements

Memory is a key parameter for the performance of Exchange server roles. This section provides design guidelines for allocating guest memory for VMs pertaining to various Exchange server roles in a virtualized environment, based on the guidelines set by Microsoft. Memory required for an Exchange server VM is dictated by the particular server role. However, when the server roles execute in a virtualized environment, the memory requirements for the hypervisor and the guest OS should also be considered in addition to the memory requirements for the specific Exchange server role(s). A set of standard guidelines pertaining to host and guest memory configuration, prescribed by VMware, is summarized here:

- The memory for the ESXi hosts running Exchange workloads should not be over-committed. This rule can be implemented by configuring the memory reservations per VM or locking the memory with the VM.

NOTE: The memory reservations should be configured carefully because they may have impacts on VMware VMotion™. A VM is only capable of VMotion if the target ESXi host has free physical memory equal to or greater than the size of the reservation.

- Configuring the memory reservations for the VM results in improved performance, as it eliminates hypervisor swapping. Hypervisor swapping is typically the last option for reclaiming the memory from the VMs and has negative impact on the VM performance.

²Introduction to Exchange Server Sizing:

<http://i.dell.com/sites/content/business/solutions/whitepapers/en/Documents/exchange-server-2010-sizing.pdf>

- The memory for the Exchange VMs should be rightly sized and the reservations and limits should be set appropriately. With reservations set, the memory will not be utilized optimally if the Exchange VMs are not utilizing the configured memory.
- The balloon driver should not be disabled. This driver gets installed with VMware Tools, implying that the VMware tools should be installed on all the VMs running Exchange server.
- Although obvious, it should be noted that the host memory size should be greater than the guest memory size, especially during the VMotion.

Memory Requirements for Mailbox Server

In order to identify memory requirements for mailbox server role, firstly, the database cache size is determined for the given user profile based on a lookup table provided by Microsoft. The Extensible Storage Engine (ESE) uses a database cache to reduce I/O operations. In general, the more database cache available, the less I/O generated on an Exchange Server 2010 mailbox server. However, there's a point at which adding database caches no longer results in a significant reduction in IOPS. Details for identifying the memory requirements for mailbox server role can be found in Appendix D: Sizing Memory for Mailbox Server Role.

Memory Requirements for the Hub Transport and Client Access Server Roles

The guidelines for identifying the memory requirements are quite straightforward, once the number of CPU cores has been identified for Hub Transport and Client Access server roles. Microsoft recommends a minimum of 4GB memory for the VMs running Hub transport and Client Access server roles. The recommendations are described below. Apart from the server roles, the hypervisor and guest memory requirements should be taken into consideration.

Table 3. Recommended Memory Configurations for Hub Transport and Client Access

Exchange server roles	Minimum supported	Recommended maximum per core
Hub Transport server role	4GB	1GB
Client Access server role	4GB	2GB
Client Access and Hub Transport combined server role (Client Access and Hub Transport server roles running on the same physical server)	4GB	2GB

The Design framework discussed thus far prescribes a structured approach toward identifying the virtual resource configuration necessary for an Exchange solution. This approach does not tie the resource configuration to the underlying hardware resources. As long as the virtual resource requirements are satisfied by underlying infrastructure, the solution can be implemented. The design framework is put to use to explain a reference implementation of an example Exchange solution with Active System 800v.

Reference Implementation for 5,000 Mailboxes on Active System 800v

The reference implementation discussed in this guide uses Active System 800v configuration with eight ESXi hosts as described in Table 1. However, the entire compute capacity available with eight ESXi hosts configuration was not used for deploying this Exchange reference implementation. The virtual resources required were sized for 5,000 mailboxes and the solution requirements shown in Table 4 govern the number of physical ESXi hosts used for Exchange virtual infrastructure deployment. In this reference implementation, three EqualLogic PS6110X arrays were used - two arrays for hosting the Exchange 2010 databases and one for the virtual machine store. The details on identifying the number of arrays for Exchange databases are further explained in Capacity Requirements per Copy. This indicates the minimum required storage capacity for implementation of 5,000 mailboxes, described in this paper.

Table 4. Example Solution Requirements for Exchange Server

Number of Mailboxes	5,000
Average user I/O profile (messages/day)	0.18 IOPS (~ 150 messages per day)
Average mailbox size limit	Up to 1024 MB
Number of DAGs	1
Total active/passive copies per database	2 (1 active + 1 passive)

The reference implementation takes into consideration the virtualization and application best practices discussed in the design framework to identify the resource requirements and implement the solution. To apply the design concepts, example Exchange solution requirements, shown in Table 4, have been considered. The implementation mainly focuses on Mailbox, Hub Transport and Client Access server roles. The Unified Messaging and Edge Transport server roles are not covered as part of the implementation. Their implementation is left to readers' discretion for scenarios specific to their production environment. It is assumed that the Active Directory Domain controller already exists in the environment; its implementation is not in the scope of this paper. The implementation begins with sizing the storage and then goes through the process of identifying the network, CPU and memory requirements for a purely virtualized environment.

A quick summary of infrastructure services provided by Active System 800v that are useful for implementing the solution is included here.

- Configuration and creation of application specific data LUNs on the shared storage.
- Configuration and creation of volumes for guest OS image on the shared storage.
- Ability to ensure bandwidth guarantees for application networks.
- Configuration of vSwitches and virtual adapters with Exchange-specific VLANs.
- Creation of VMs with specific settings for resource utilization.
- Reserving and limiting resources and not oversubscribing them for individual VMs.
- Implementing resource allocation and scheduling rules.
- Ability to move VMs within a cluster of hypervisors.

Sizing and Configuring Storage

This section describes storage sizing and configuration for the example solution requirements shown in Table 4. Active System 800v provides ample amounts of iSCSI storage, with up to eight EqualLogic PS6110 storage arrays. For more information on storage options provided by Active System 800v, please refer to the implementation guide provided with Active System 800v documentation. Based on pilot tests with Exchange Jetstress tool, RAID-50 configuration was chosen for Exchange database storage with EqualLogic PS6110X arrays consisting of 900GB 10K SAS drives.

Capacity Requirements per Copy

Exchange Server is a storage-intensive application. The capacity sizing for Exchange depends on a number of factors, such as the retention policy, dumpster space, database search index, database overhead, and space for transactional log and 20% estimated free space. For the example, the mailbox profile described in Table 4. A 1GB mailbox size with a 150 messages/day profile with 14 days retention window and average message size of 75KB requires 11MB of whitespace and 197MB of Dumpster, which makes the mailbox size on disk 1232 MB. Considering 20% data overhead, the total database capacity required for 5,000 mailboxes can be calculated as: $(5,000 \times 1232 \times 1.2) / 1024 \text{ GB} = 7218.75 \text{ GB}$. As per the Microsoft recommendations, the database index size is considered as 10% of the total database capacity required, which equals 721.875 GB in our example. The total database and index capacity required is 7940.625 GB. The total volume capacity required for the database and the index is calculated by taking into consideration 20% volume free space. Thus, the volume capacity required for the database and index combined is 9925.781 GB. The last consideration is the capacity required for log files. For the given mailbox profile, the required log space can be calculated to be 489.45 GB. Considering 20% volume free space for logs, the total volume capacity required for logs is 612.43GB.

The total volume capacity required for a single copy of our example Exchange database can be calculated by adding the total volume capacity for database and index and the total volume capacity for logs; which equals $9925.781\text{GB} + 612.43\text{GB} = 10538.2185\text{GB}$, or approximately 10.6 TB. For an EqualLogic PS6110X array with 900 GB drives storage enclosure, usable space (RAID 50) = 14.48 TB, which indicates that one array / member is sufficient for one copy of the Exchange database. A two-copy DAG will require two such arrays.

5,000 mailboxes were planned to be deployed over 9 databases, implying approximately 556 mailboxes per database, with each database having capacity of 1.18 TB. Additional 1.18 TB for restore LUN were reserved. Three mailbox server VMs would be needed in order to evenly distribute the derived number of databases. A standalone installation of a mailbox server role as compared to multi-role installation would be efficient. The Hub Transport and the Client Access server roles can be collocated on a single virtual server for this example scenario. Separating the guest OS images from Exchange databases is recommended. Hence, an additional EqualLogic storage array is needed to store the VMs required for Exchange server.

IOPS Requirement per Copy

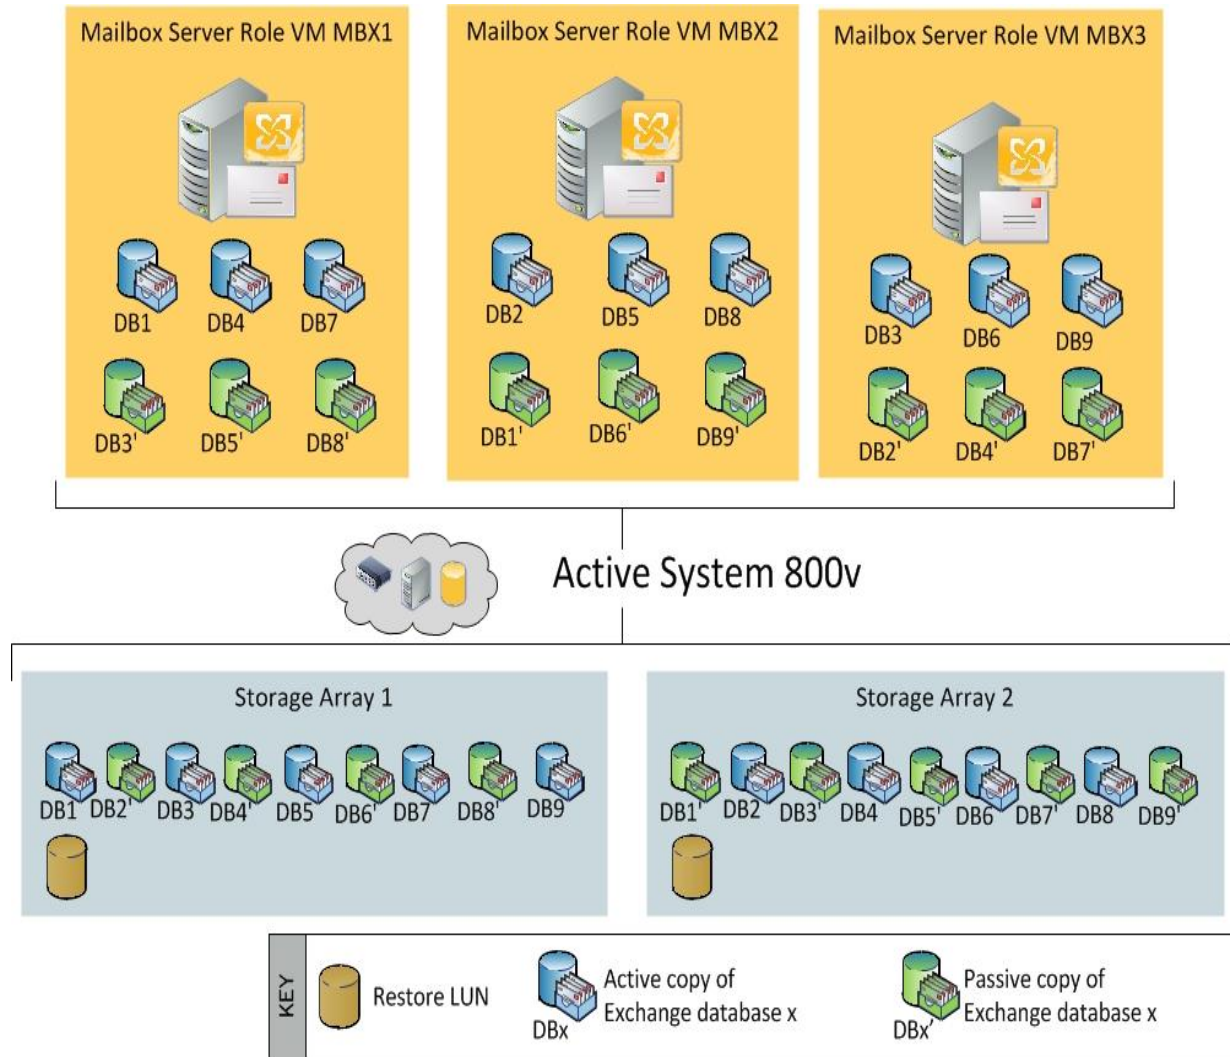
The target I/O profile in the example solution requires 150 messages a day per active mailbox copy = 0.15 IOPs / mailbox. Considering an I/O overhead of 20%, about 0.18 IOPS per mailbox are to be considered. Thus, for a target of 5,000 mailboxes, a total of $(0.18 * 5,000) = 900$ IOPs are required.

For a detailed approach on Exchange sizing, please refer to ‘Mailbox role storage capacity calculations’ and ‘Mailbox role IOPS calculations’ sections in the Introduction to Microsoft Exchange Server 2010 sizing³ paper.

DAG Storage Configuration

Considering the guidelines for identifying the storage requirements, two EqualLogic PS 6110X members were utilized to host the two-copy DAG described in the example solution. A logical layout of the DAG design involving active and passive database copies is shown in Figure 2.

Figure 2. A Logical Layout for the DAG Storage



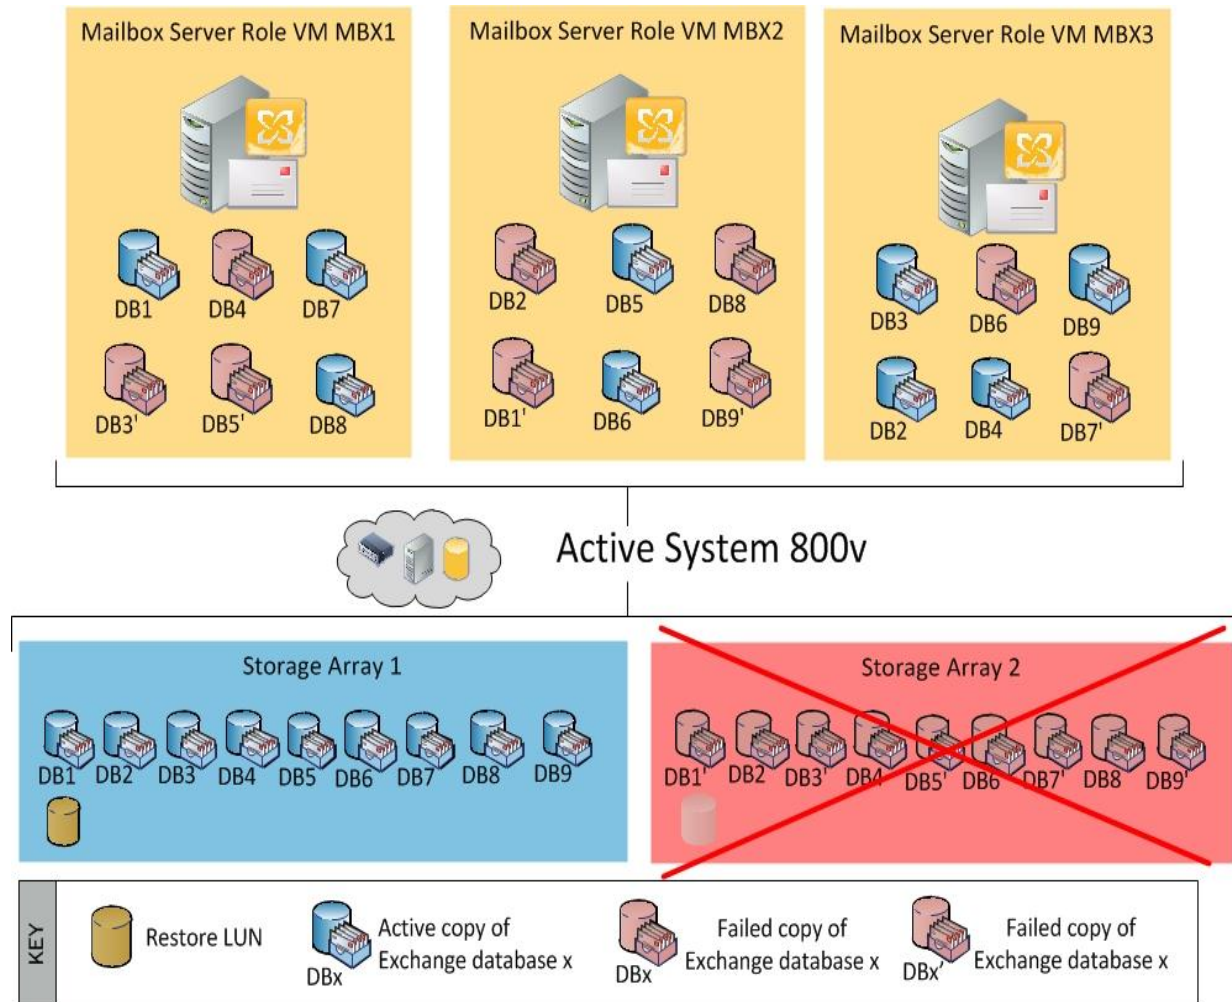
³Introduction to Exchange Server Sizing:

<http://i.dell.com/sites/content/business/solutions/whitepapers/en/Documents/exchange-server-2010-sizing.pdf>

Nine volumes per EqualLogic member were created to account for the nine databases distributed over a two-copy DAG. Each Exchange database copy is mapped to one volume. Allocation of the active and passive databases over the storage arrays and the mailbox VMs is planned in such a way that either the storage or the server failure will result in fairly even distribution of databases across the VMs running the mailbox server role. Snapshots were disabled for the database volumes in order to utilize maximum storage space. Access to these volumes was based on the respective iSCSI subnets instead of iSCSI initiator names. Multiple connections to these volumes were enabled for storage multi-path accesses. The Figures below show impact on the DAG in case of storage and server failure. These failures are considered mutually exclusive.

In order to handle a storage array failure, it is recommended that the active and passive copy of a database not be stored on the same storage array. Assume the storage array 2 encounters a failure. The active copies of the databases hosted on storage array 2, namely DB2, DB4, DB6, and DB8 encounter failure. And the passive copies of DB3, DB5, DB7, and DB9 become unavailable. As a consequence, the respective passive copies of the databases DB2, DB4, DB6, and DB8 hosted on storage 1 get activated and continue to provide the email service. The failure scenario is shown in Figure 3.

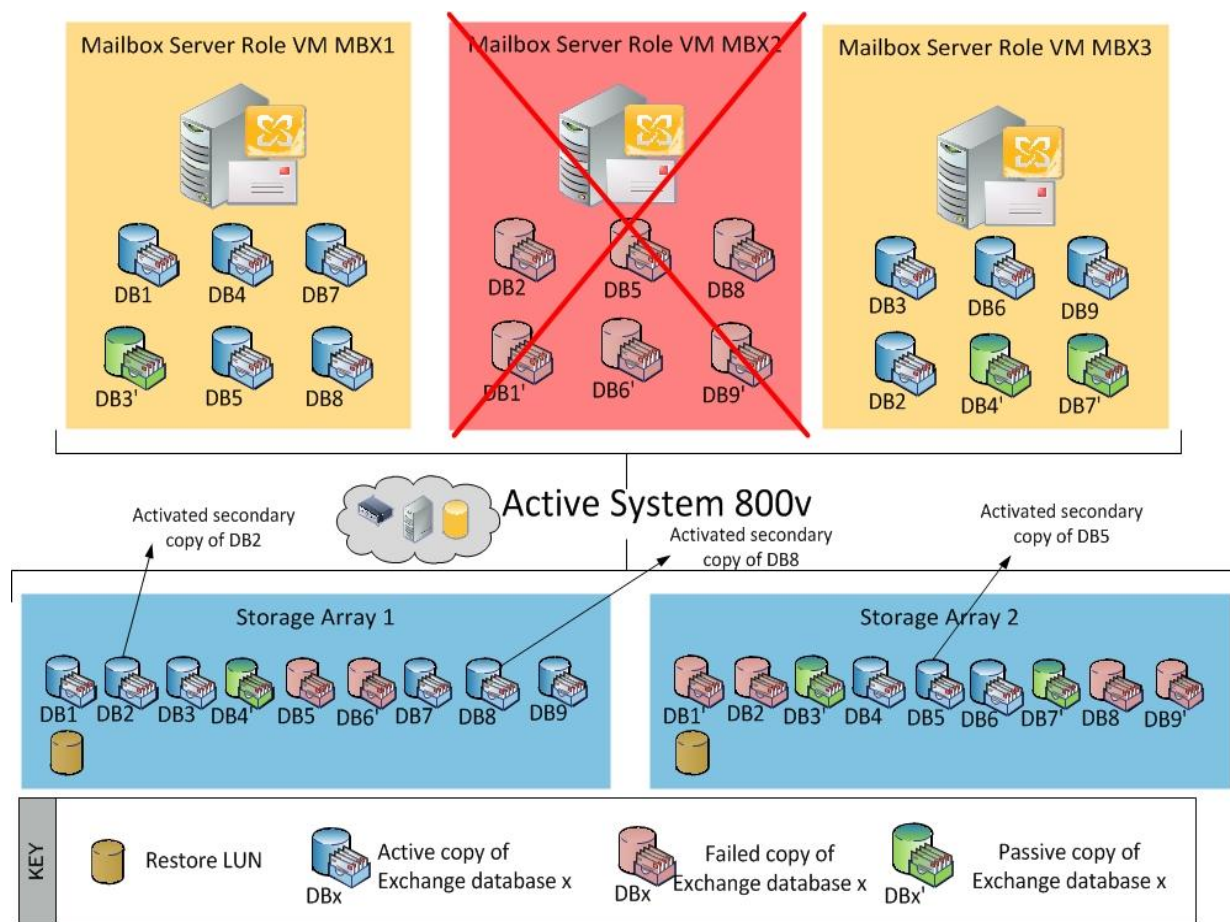
Figure 3. DAG State After a Storage Array Failure



Based on the database distribution, Storage Array 1 helps to provide email service in the absence of the other storage array. This implies that, from a sizing verification perspective, it is necessary to ensure that one storage array is sufficient to meet the peak load requirements of the given mail profile.

Another failure scenario is a hypervisor host failure. As it is ensured that no more than one mailbox server VM will be running on single hypervisor host, one of the mailbox server VMs hosted on that hypervisor will experience downtime when the hypervisor host fails. Let us assume that the hypervisor hosting MBX2 encounters failure. The DAG experiences a node failure, while the VMware High Availability services redistribute all the VMs on the failed hypervisor on remaining hosts in the cluster. When MBX2 goes down, the active database copies mounted on MBX2, namely DB2, DB5, and DB8, experience failure. Exchange native HA activates the respective passive copies of these databases on MBX1 and MBX3 and continues to provide the email service. The scenario is shown in Figure 4.

Figure 4. DAG State After Server Failure



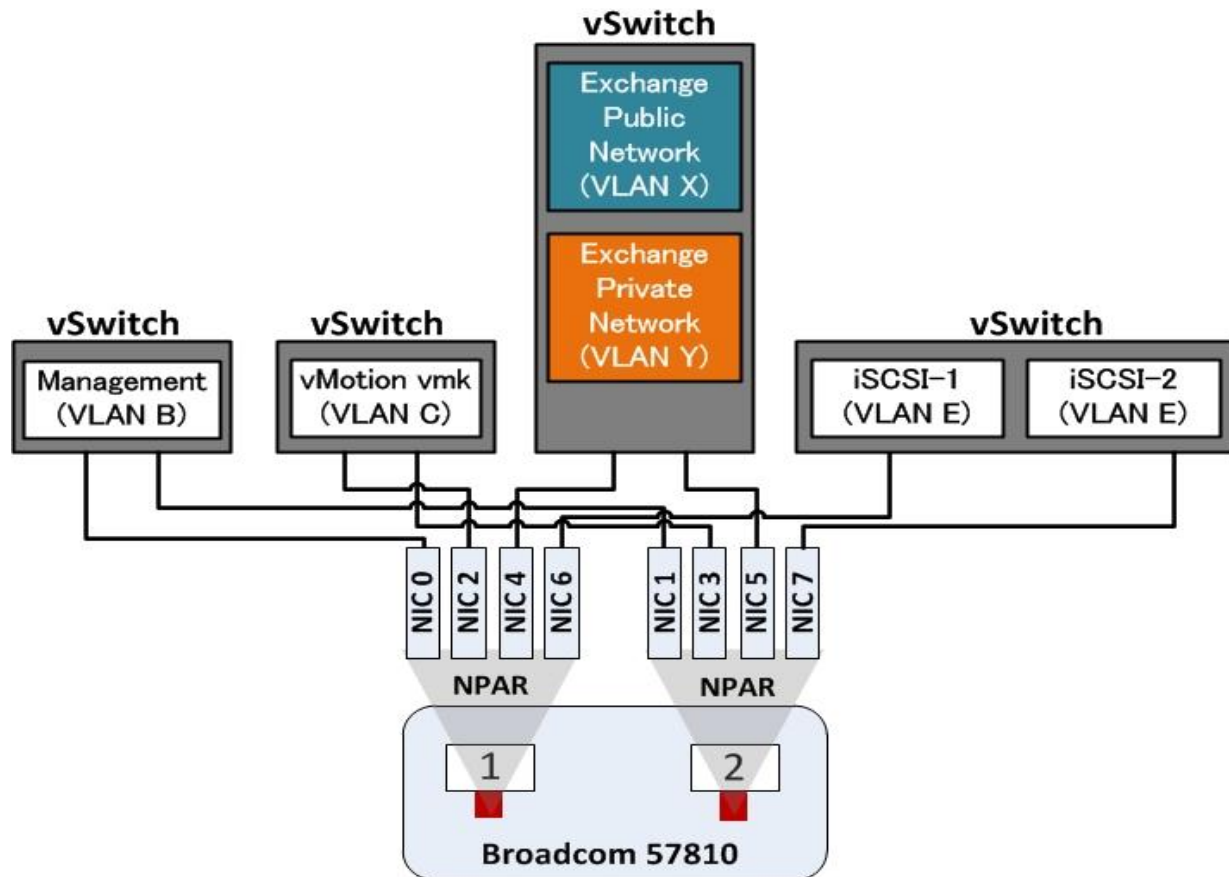
Identifying and Configuring Network Requirements

Active System 800v ensures network virtualization best practices and provides infrastructure services to implement the application networking best practices. This section describes the implementation of Exchange-specific networks in Active System 800v environment. It addresses two aspects of network configurations, namely the host networking and the VM networking.

Host Network Configurations

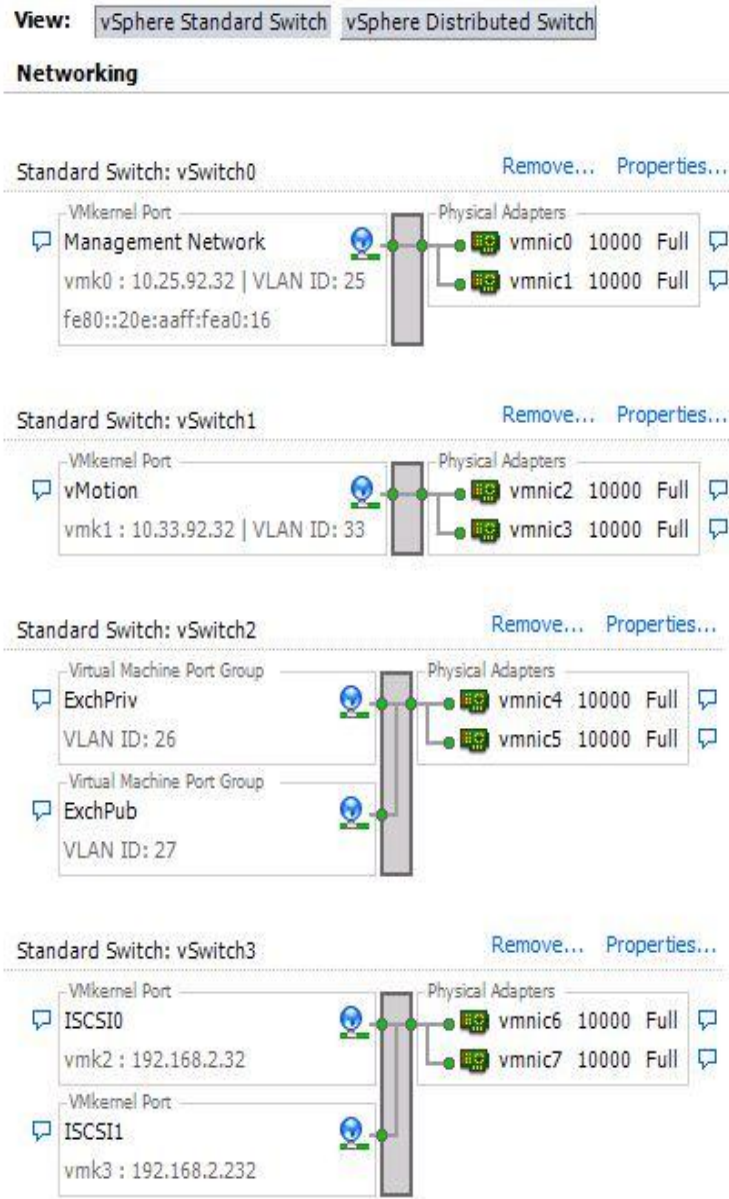
MAPI and the DAG replication also referred to as Exchange Public and Private Networks respectively, are the two primary networks Exchange server deployment relies on. This means at least two VLANs are needed, corresponding to the MAPI and DAG replication network. An additional VLAN is needed for the iSCSI EqualLogic storage. Two separate virtual switches are planned to be used: one for Exchange Public (MAPI) and Private (DAG replication) networks and one for the iSCSI storage. It should be ensured that the network configurations must be consistent across all the hosts in a cluster. This ensures the connectivity to the required storage LUNs, irrespective of the VM placements on the hosts in the cluster. At the host level, a 10Gbps dual port Broadcom 57810 network adapter interfaces with the converged network infrastructure. For better utilization of the bandwidth and resiliency at the network adapter, hardware partitioning of the network adapter is provided by Active System 800v. The two available 10Gbps ports on the Broadcom 57810 are divided into eight hardware partition using the NPAR. The iSCSI offload is enabled on two of the partitions. Four different networks, namely Management, vMotion, Exchange public and private, and iSCSI, map to these eight hardware partitions, as shown in the Figure 5. Although there is no need to configure the NPAR configurations, it is recommended to understand the mapping of the Exchange application networks to the physical network adapters. For purposes of the reference implementation, VLAN X is set to 27 and VLAN Y is set to 26. For details on configuring VLANs on Active System 800v, please refer to the Active System 800v documentation.

Figure 5. Active System 800v NPAR Configuration for Exchange Server



A sample implementation of host networking for Exchange on Active System 800 is shown in a screen capture in Figure 6. This includes the default host networking provided by Active System 800 and Exchange-specific host networks. Virtual Switch 2 is configured to have Exchange Public (MAPI) and private (DAG replication) networks. These two networks are identified by two different VLAN IDs, namely 27 and 26 respectively.

Figure 6. Sample Implementation of Host Network Configuration



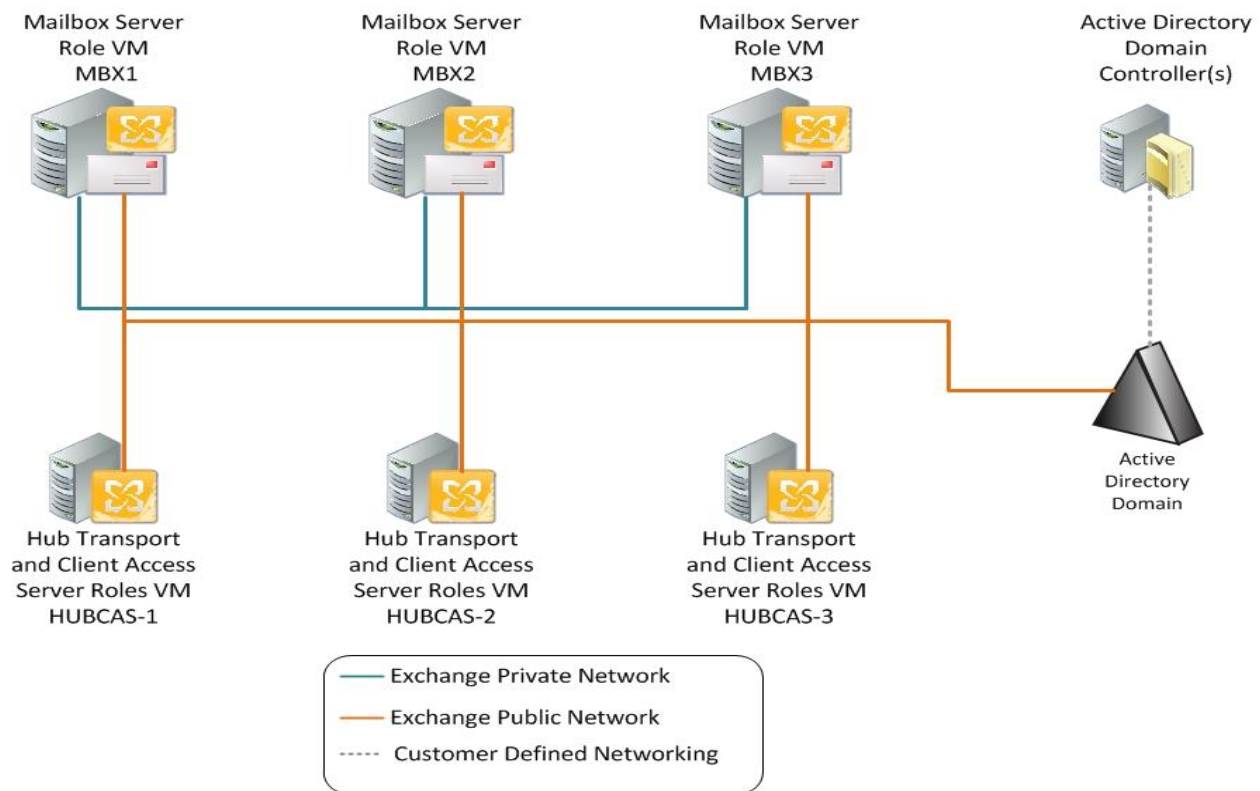
Active System 800 provides a preconfigured network configuration for iSCSI storage. Once the host networking is configured, the storage LUNs pertaining to the Exchange databases and the VM OS image must be discovered for each virtual HBA for storage on each host in the cluster.

VM Network Configurations

The VMs running mailbox server role need the Exchange Public and the Exchange Private networks. These two networks should be configured on separate network adapters, implying at least two virtual network adapters are needed for mailbox server VMs. A VM running Hub Transport and Client Access server collocated server roles (HUBCAS) needs to communicate with the Active Directory through MAPI network, implying that the VM needs at least one virtual network adapter. The logical allocation of virtual network adapters is shown in Figure 7.

For this reference implementation, it is assumed that the Active Directory Domain Controller already exists at the customer site and the implementation of networking related to the Active Directory is left to the customer's discretion.

Figure 7. VM Logical Networking



This section summarized the network configuration aspects for the reference implementation from a host and VM perspective. These network implementation guidelines are independent of the number of mailbox users in the solution.

Identifying and Configuring CPU Requirements

This section describes identifying the CPU requirements primarily for the Mailbox, the Hub transport, and the Client Access server roles. In this example scenario, mailbox server roles are kept standalone while the Hub transport and the Client Access server roles are collocated from an installation perspective.

In a normal operational DAG, each mailbox server VM hosts three active and three passive databases, resulting into 1667 active and 1667 passive mailboxes. However, in a failure scenario (either one storage member or one Mailbox server VM fails), with two mailbox server VMs operational, in worst case scenario, one of the VM will have to host five active and one passive database, resulting in a maximum of 2,780 active and 560 passive mailboxes. Assuming this worst case scenario, we size four vCPUs per Mailbox VM.

As described in the 'Virtual CPU Requirements for the Hub Transport and Client Access Server Roles' section, the recommended processor core ratio for Mailbox server role to (Hub Transport + Client Access) server roles is 1:1. The planned vCPUs per mailbox VM is 4; hence we plan to configure four vCPUs per VM running collocated Hub transport and Client Access server roles (HUBCAS). While configuring the VM with multiple vCPUs, the reader can either use one socket and multiple cores per socket or multiple sockets and one core per socket. This reference implementation used a single virtual socket and multiple cores per socket.

Identifying and Configuring Memory Requirements

The example in consideration requires a total of 150 messages to be sent and received per mailbox per day. Based on the message profile, 9MB of database cache is required as per the database cache requirements described in the 'Memory Requirements for Mailbox Server' section. Based on the DAG allocation shown earlier, with a normal operational DAG, each mailbox server will host $5,000/3 = 1,667$ mailboxes; however, in a failover scenario, a mailbox server VM may end up hosting up to five databases, or up to $5 * 556 = 2,780$ active mailboxes.

Total database cache required per server = Number of active mailboxes per server x average database cache per Mailbox

Total database cache required per server = $2,780 \times 9 = 25\text{GB}$

As described in Memory Requirements for Mailbox Server section, memory required per mailbox server VM is at least 32GB. An additional memory of 4GB should be allocated for the guest OS. For collocated Hub transport and Client Access servers, at least 2GB of memory per core is recommended. Because 4 cores are recommended per HUBCAS VM, at least 8GB of memory are essential per HUBCAS VM. To account for additional memory requirements of the guest OS, an additional 4GB memory must be provided. Thus the mailbox server VMs were configured with 36GB of memory, while the HUBCAS VMs were configured with 12GB of memory.

Configuring VM Attributes

The resource requirements in terms of CPU, memory, networking, and storage derived based on the design guidelines so far are summarized in Table 5. This section describes the configuration of VMs specific to Exchange server roles. It is recommended to use the VM templates provisioned with the OS and Exchange server roles prerequisites to install and deploy Exchange server roles. The recommended OS for Exchange server 2010 SP2 is Microsoft Windows Server® 2008 R2.

Table 5. Resource Requirements for Exchange Server Solution

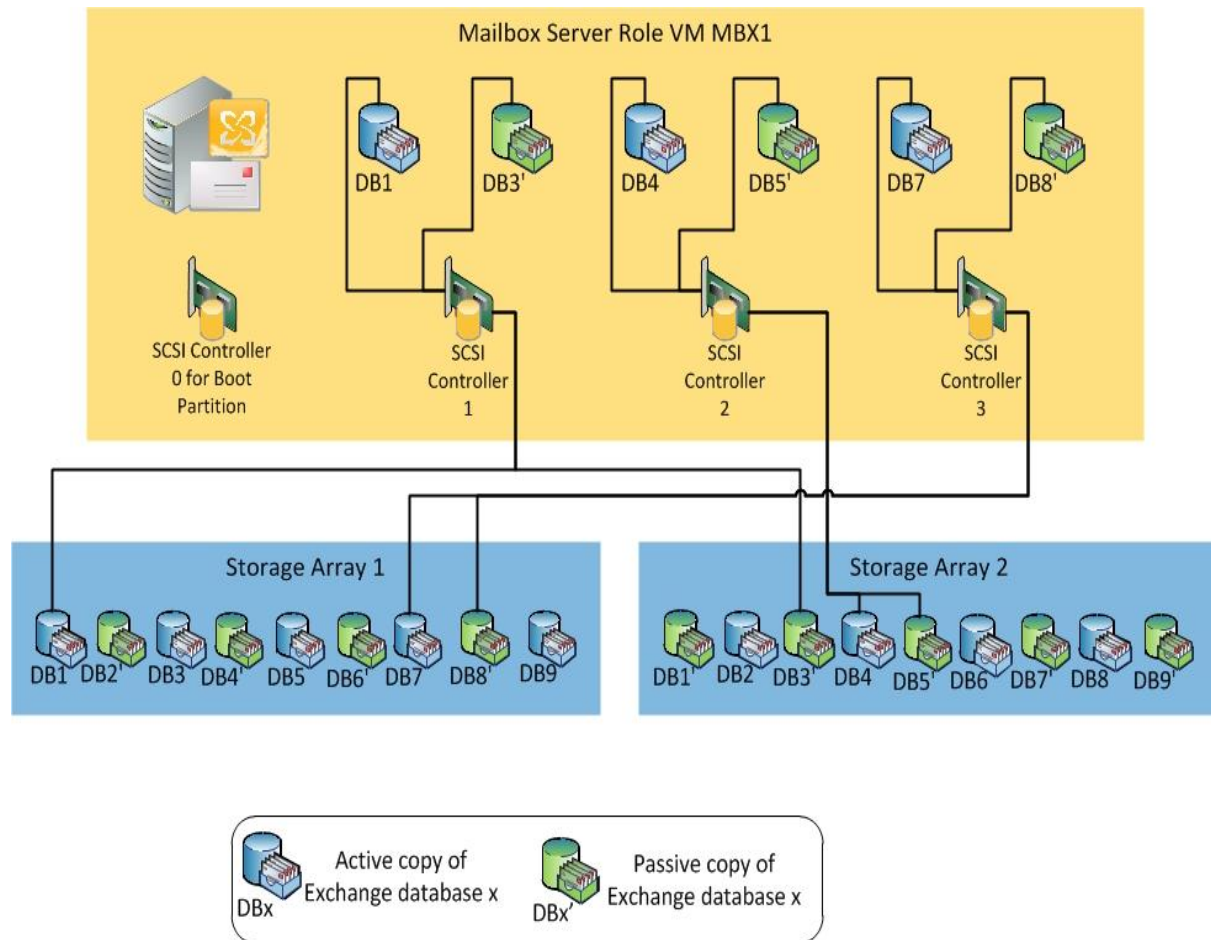
	Mailbox server role	Hub Transport and Client Access server role collocated
Number of VMs	3	3
Virtual processors (vCPUs)	4 (1 socket, 4 cores/socket)	4 (1 socket, 4 cores/socket)
Memory (GB)	36	12
Network virtual adapters	2x VMXNET3 vNIC	1x VMXNET3 vNIC
Storage Controller	LSI Logic SAS	LSI Logic SAS
Storage for Guest OS image	150GB Thick Provisioned lazy zeroed VMDK	150GB Thick Provisioned lazy zeroed VMDK

Exchange server roles require dedicated resources for memory and processing resources. Resource allocations can be controlled using VM attributes like Shares, Reservations, and Limit.

- Shares enable a priority for VMs relative to other VMs in their hierarchy. Shares are typically specified in the VM settings as high, normal, or low, corresponding to the ratio 4:2:1. If a VM has twice as many shares for a resource as another VM, it is entitled to consume twice as much of that resource. Custom shares can also be assigned to express a proportional weight to each VM. Specifying shares is most applicable only with regard to VMs at the same hierarchy in a resource pool tree.
- Reservations help guarantee minimum resource requirements for CPU and memory. A VM, when attempting to power on, will start only when the reservation values for the resources set are satisfied. Also, the reservation values are guaranteed under peak workload.
- Limits help in specifying the upper bound for CPU and memory. Limits imply that the resource allocation for a VM will never exceed the values specified as Limits.

It is recommended to set reservations to the derived CPU and memory requirements, in order to guarantee the required resources to VMs. For memory resources, one of the options is to lock the memory for a particular VM. This ensures that the memory consumed by and granted to a VM always remains equal. There are multiple ways to introduce the storage LUNs, pertaining to the Exchange database, to the VM running a mailbox server role. Based on the design guidelines, it was decided to map the Exchange database storage LUNs to the mailbox server role VMs using RDM volumes. Figure 8 depicts a logical allocation of the RDM disks to the virtual SCSI controllers.

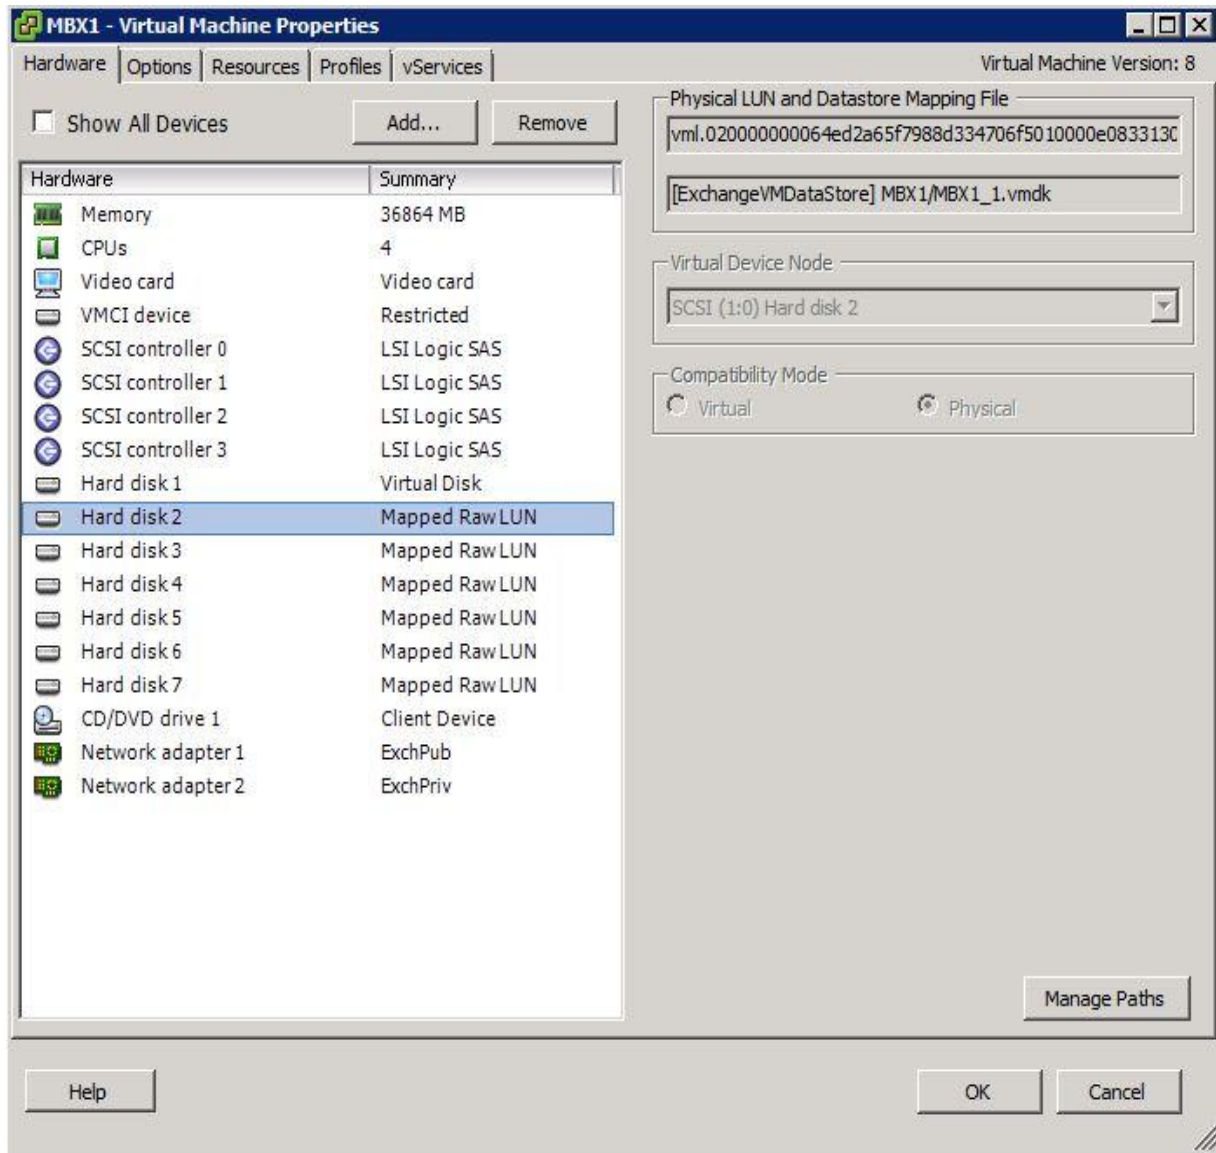
Figure 8. Allocating RDM Disks to Virtual SCSI Controllers



It is vital to accurately balance the storage LUNs across the available virtual SCSI adapters. VMware provides four virtual SCSI adapters for each VM and fifteen disks per virtual SCSI adapter. A virtual device node can be mapped as “SCSI (controller: disk)”. For the reference implementation under consideration, each mailbox VM has six database storage LUNs and an additional LUN for the guest OS image. Typically, controller 0 is utilized to map the device node for the guest OS image. It is not recommended to use the same SCSI controller as that of the guest OS image for Exchange database LUNs. The remaining three SCSI controllers were used to map six device nodes as six database LUNs. Out of the six LUNs, three LUNs correspond to active while the remaining three correspond to passive databases. Each of the three available SCSI controllers is mapped to an equal number of active and passive database LUNs, resulting in one active and one passive database LUN per SCSI controller. Once the database storage LUNs are mapped, the discovered raw storage volumes are required to be activated and initialized with MBR or GPT and configured as NTFS volumes with 64K Allocation unit size.

An example VM configuration with CPU, memory, network adapters and Exchange database LUNs for mailbox server role is shown in Figure 9.

Figure 9. VM Attributes for Mailbox Server Role



NOTE: It is recommended to use Network Load Balancing (NLB) to balance the load across Client Access servers.

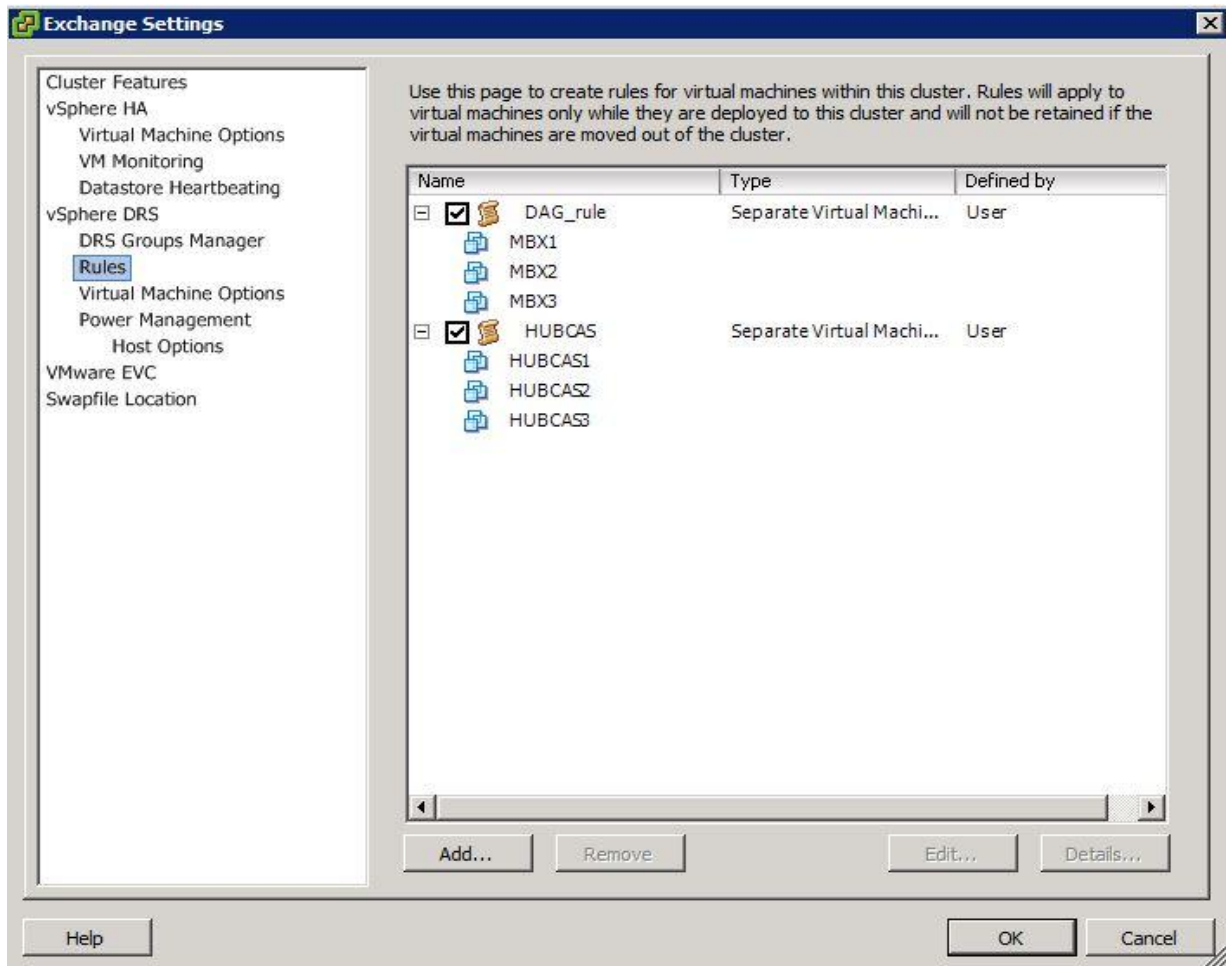
VM Placements

The derived virtual resource requirements demand at least three VMware ESXi hosts for a steady state DAG, with all the DAG nodes and HUBCAS VMs functioning. However, the DAG with one node and one HUBCAS VM failure can operate on two ESXi hosts capable of satisfying the resource requirements of operational Exchange server VMs. The reference solution requires at least two mailbox and two HUBCAS VMs operational at any given time. In order to provide tolerance to one of the host's failure, separate VM rules were implemented for both the mailbox and HUBCAS server VMs.

1. Mailbox servers must reside on separate ESXi hosts in any scenario from a DAG operation perspective.
2. HUBCAS servers must reside on separate ESXi hosts in any scenario from a performance perspective.

Figure 10 shows the DRS rules for the Exchange cluster, used in the reference implementation. It should be noted that as long as the virtual resource requirements and the placement policies of the Exchange server VMs are satisfied, they can be placed to run with VMs running other workloads.

Figure 10. VM Placements



Solution Summary

Considering the infrastructure and the application perspective, a solution summary for the requirements considered in the reference implementation is shown here.

Table 6. Solution Summary

Exchange Server 2010 Major Components	
Hub Transport and Client Access VMs	
Number of VMs	3
Guest Operating System	Windows Server 2008 R2 SP1
VM Configuration	4 x vCPUs; 12 GB Virtual Memory per VM
Networks	1 x VMXNET3 for Exchange Public Network per VM
Software	Exchange Server 2010 SP1
Mailbox Server VMs	
Number of VMs	3
Guest Operating System	Windows Server 2008 R2 SP1
VM Configuration	4 x vCPUs; 36 GB Virtual Memory per VM
Networks	2 x VMXNET3 per VM; one each for Exchange Public Network and Exchange Private Network
Software	Exchange Server 2010 SP1
Storage Volumes	
Number of Storage arrays	2 x EqualLogic PS6110X for Exchange databases
	1x EqualLogic PS6110X for VM store
RAID configuration	RAID 50
Active Mailbox Databases	9 x 1.2 TB Volumes
Passive Mailbox Databases	9 x 1.2 TB Volumes
Restore LUNs	2 x 1.2 TB Volumes

NOTE: Throughout the implementation process, it is assumed that the Active Directory Domain Controller already exists in the customer environment.

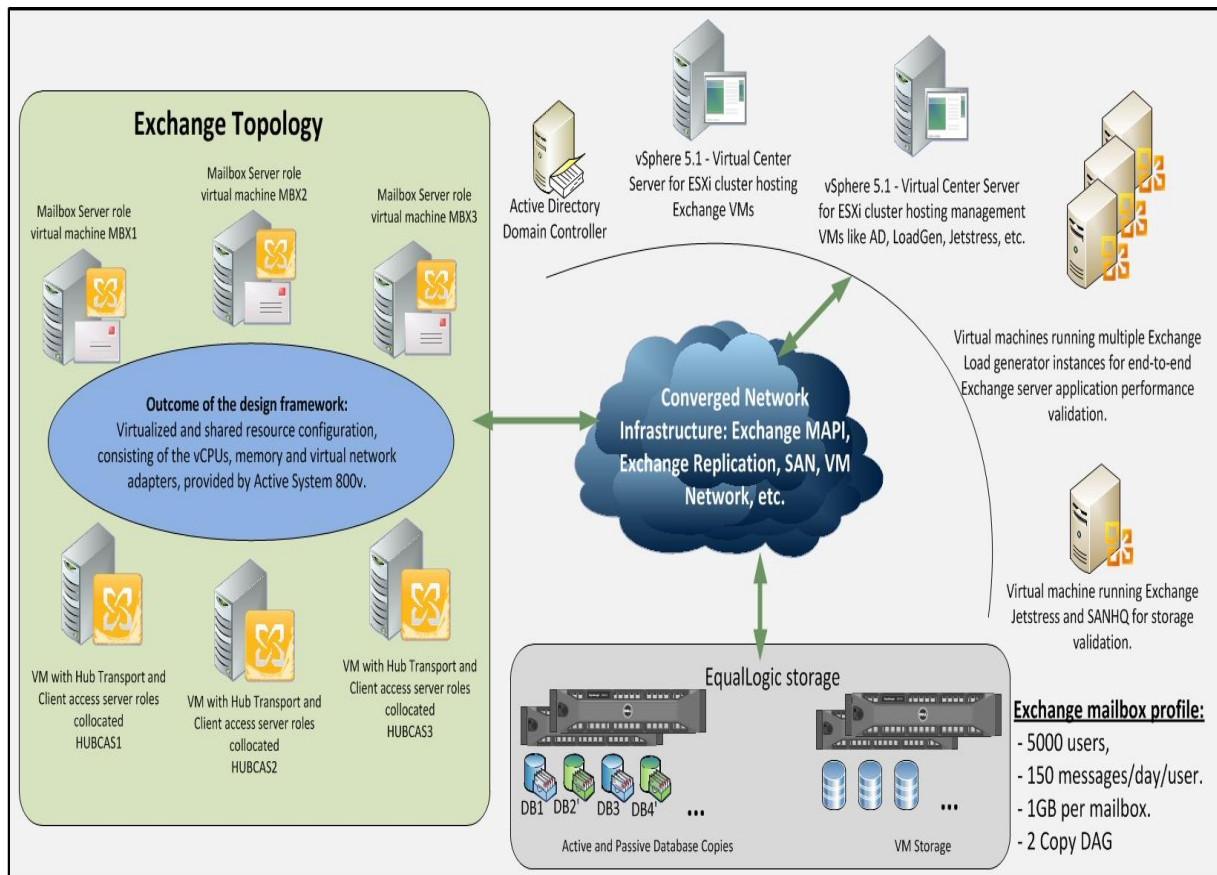
Reference Implementation Verification of 5,000 Mailbox Users

This section mainly provides proof points for the performance and resource allocation and utilization for the reference implementation of the example Exchange solution for 5,000 users on Active System 800v. The solution verification consisted of verifying the performance and DRS rules.

Performance Verification

The performance verification was broadly divided into two categories: verification of the storage subsystem performance and the end-to-end application performance verification. Microsoft recommends using standard performance verification tools Jetstress and Load Generator to ensure that the theoretical sizing of the application components meets the performance requirements. Microsoft Exchange Jetstress was used for verifying storage performance, and Microsoft Load Generator tool was used for verifying end-to-end Exchange application performance. Performance tools were configured such that they were isolated from the cluster of hosts running Exchange server roles; however, they interacted externally with the cluster. A logical representation of the test tools used in the test bed is shown in Figure 11. The performance verification considered primarily the results from Microsoft recommended tools for Exchange server.

Figure 11. Test Bed for Exchange Server Verification with Active System 800v



Verifying Storage Requirements Using Jetstress

The verification storage subsystem performance was intended to ensure that the storage sized theoretically, for the Exchange databases, meets the performance expectations for the given number of mailbox users and the mailbox profile. Running the test with Jetstress measures how well a storage system can perform and whether the storage subsystem meets the sizing requirements for a given Exchange mailbox profile. The Jetstress Disk Subsystem Throughput test was performed in order to identify how well the storage performs at peak load while staying in the latency threshold established by Microsoft Exchange. The reference implementation describes the use of two storage arrays. Assuming a storage array failure as the worst case operational scenario, the test was run for nine active databases on a single storage array. One storage array consisted of 24x 2.5' 900GB drives in a RAID-50 configuration. The ratio of Exchange database to the storage volume was kept 1:1, with each storage volume of 1.2 TB. The Jetstress test was run and the test results were collected from within the VM. The nine active databases were distributed evenly across three virtual SCSI adapters of the VM running the test.

The transactional I/O performance in terms of achieved average transactional I/O per second was observed. The overall results from the Jetstress Disk Subsystem Throughput test depicted impressive performance, implying faster Exchange performance in terms of data transfers, and significant increase in IOPS per server. The results of the Disk Subsystem Throughput test verified that the storage can handle the peak load for the mailbox profile in consideration. The mailbox profile for the reference implementation used for running Jetstress Mailbox Profile test is shown in Table 7.

Table 7. Mailbox Profile for Jetstress Test

Number of Mailboxes	Mailbox size (GB)	Target IOPS	Storage	RAID type	Number of Exchange databases	Volume size	I/O profile
5,000	1 GB	900	EqualLogic PS6110X	RAID-50	9 (1 DB per volume)	1.2 TB	150 messages a day (0.18 tested)

The Jetstress Mailbox Profile test verifies whether the storage system meets or exceeds the planned Exchange mailbox profile. The configuration shown in Table 7 was used to run the mailbox profile test with a planned target of 900 IOPS. The achieved target was 1315 IOPS. The results indicated that the storage was able to exceed the target transactional IOPS, staying well within the latency requirements.

Verifying the End-to-End Performance of Exchange Server with Active System 800v

When verifying server and storage for Exchange Server 2010, a best practice is to simulate the worst case scenario under anticipated peak. Based on a number of data sets, peak load is generally twice the average workload throughout the remainder of the work day. In order to determine whether the converged infrastructure can handle both normal as well as peak loads, the test plan consisted of running Load Generator tests with peak and normal loads. To simulate peak load, the simulation day was set to four hours and test duration was set to eight hours, and to simulate normal load, the simulation day and test duration were both set to eight hours. Based on the DAG configuration, one DAG node failure could be tolerated. The LoadGen workload was configured to simulate Outlook 2007

Online with Outlook-150 profile. This profile simulated sending and receiving 150 messages per mailbox per day.

The performance test plan intended to verify the end-to-end performance of Exchange server 2010 in Active System 800v environment with a steady state configuration (DAG with all nodes functional) as well as the failed host configuration (DAG with one failed node). Both the configurations were subjected to run under peak and normal loads. Three Load Generator instances were used to distribute the generated load across the three DAG nodes. Table 8 briefly describes the test plan for Load Generator tests.

Table 8. Performance Test Plan for Exchange Load Generator

Tests	Test description	DAG state	Database distribution
Test 1	Peak load test with steady state configuration: Simulation day: 4 hours Test length: 8 hours	Host 1: MBX1, HUBCAS1 Host 2: MBX2, HUBCAS2 Host 3: MBX3, HUBCAS3	MBX1: DB1, DB4, DB7 MBX2: DB2, DB5, DB8 MBX3: DB3, DB6, DB9
Test 2	Normal load test with steady state configuration: Simulation day: 8 hours Test length: 8 hours	Host 1: MBX1, HUBCAS1 Host 2: MBX2, HUBCAS2 Host 3: MBX3, HUBCAS3	MBX1: DB1, DB4, DB7 MBX2: DB2, DB5, DB8 MBX3: DB3, DB6, DB9
Test 3	Peak load test with failed host configuration: Simulation day: 4 hours Test length: 8 hours	Host 1: MBX1, HUBCAS1 Host 2: MBX2, HUBCAS2 Host 3: FAILED	MBX1: DB1, DB4, DB7, DB3 MBX2: DB2, DB5, DB8, DB6, DB9
Test 4	Normal load test with failed host configuration: Simulation day: 8 hours Test length: 8 hours	Host 1: MBX1, HUBCAS1 Host 2: MBX2, HUBCAS2 Host 3: FAILED	MBX1: DB1, DB4, DB7, DB3 MBX2: DB2, DB5, DB8, DB6, DB9

The performance results pertaining to the middle four hours of an eight hour test are considered in the analysis. Except for the processor and memory utilization, all the application specific counters are measured within the guest OS. The processor and memory utilization has been measured from the performance reporting tool provided by vCenter.

Table 9. Test Results for VMs Running Mailbox Server Role in a Steady State Configuration

Performance Metrics	Thresholds	Test 1 (Peak load)			Test 2 (Normal load)		
		MBX3	MBX2	MBX1	MBX3	MBX2	MBX1
%Processor Utilization	< 70 %	52.952	52.244	52.644	18.009	19.602	20.237
%Memory Utilization (GB)		21.27	20.84	20.40	14.410	14.787	14.731
I/O Database Reads (Attached) Average Latency	< 20 ms	11.038	9.965	10.028	6.341	7.572	5.978
I/O Database Reads (Recovery) Average Latency	< 200 ms	16.809	10.051	10.913	25.843	5.959	14.583
I/O Database Writes (Attached) Average Latency	< 20 ms	3.411	2.158	2.390	1.256	1.354	1.334
I/O Database Writes (Recovery) Average Latency	< 200 ms	4.014	3.053	3.036	1.484	1.627	1.659
I/O Log Reads Average Latency	< 20 ms	2.541	1.757	2.090	0.920	0.972	0.967
I/O Log Writes Average Latency	< 20 ms	1.285	1.199	1.214	0.639	0.710	0.631
RPC Requests	< 70	1.925	1.323	1.898	0.725	0.855	0.723
RPC Averaged Latency	< 10 ms	1.534	1.621	1.525	1.688	1.760	1.763

The results shown in Table 9 indicate that under increased load, the Exchange performance counters pertaining to the mailbox server roles for the steady state configuration stay within the thresholds specified by Microsoft, under normal as well as peak loads. The read and write latencies pertaining to the databases and logs are well within the boundaries of the performance, indicating that the DAG configuration is rightly configured with the virtual resources, for a normal operational scenario. Table 10 captures results for HUBCAS VMs in the steady state configuration. The results shown for the Hub Transport and the Client Access servers indicate that the given design is capable of handling normal as well as peak messaging load in a steady state configuration with three HUBCAS VMs.

Table 10. Test Results for VMs Running HUBCAS Server Roles in a Steady State Configuration

Performance Metrics	Thresholds	Test 1 (Peak load)			Test 2 (Normal load)		
		HUBCAS3	HUBCAS2	HUBCAS1	HUBCAS3	HUBCAS2	HUBCAS1
%Processor Utilization	< 70 %	22.158	22.352	22.136	11.399	12.015	12.157
%Memory utilization (GB)		18.433	18.125	18.001	15.138	15.971	16.344
RPC Averaged Latency	< 250 ms	5.753	5.321	5.892	4.913	5.102	5.189
RPC Requests	< 40	1.907	1.441	1.978	0.768	0.882	0.941
Aggregate Delivery	<3000	3.349	3.318	3.386	0.786	0.164	0.786
Active Mailbox Delivery Queue Length	<250	3.154	3.145	3.173	0.752	0.161	0.752
Retry Mailbox Delivery Queue Length	<100	0.194	0.173	0.213	0.034	0.003	0.034

In order to test a degraded DAG configuration with one DAG node failure, a host failure scenario had to be considered in the test plan. To verify whether the degraded DAG can handle normal load as well as peak load, Test 3 and Test 4 were executed. Test 3 was configured for peak load while Test 4 was configured for normal load. The results pertaining to the VMs running mailbox server roles for these tests are shown in Table 11.

Table 11. Test Results for VMs Running Mailbox Server Role in a Failed Host Configuration

Performance Metrics	Thresholds	Test 3 (Peak load)		Test 4 (Normal load)	
		MBX1	MBX2	MBX1	MBX2
%Processor Utilization	< 80 %	60.071	73.847	22.353	26.854
%Memory Utilization		23.71	28.83	15.613	18.558
I/O Database Reads (Attached) Average Latency	< 20 ms	13.476	13.650	8.134	8.712
I/O Database Reads (Recovery) Average Latency	< 200 ms	12.612	56.321	13.476	53.077
I/O Database Writes (Attached) Average Latency	< 20 ms	3.306	3.013	1.280	1.693
I/O Database Writes (Recovery) Average Latency	< 200 ms	4.822	3.809	1.917	1.953
I/O Log Reads Average Latency	< 20 ms	1.550	2.421	0.979	1.109
I/O Log Writes Average Latency	< 20 ms	1.166	1.335	0.626	0.776
RPC Requests	< 70	2.205	3.020	0.925	1.506
RPC Averaged Latency	< 10 ms	1.551	1.448	1.830	2.025

For Test 3, an increase in CPU and memory utilization was observed when the DAG is in a degraded state with only two operational DAG nodes, as compared to the steady state DAG with three operational DAG nodes. Thus, it is recommended that the DAG should not remain in the degraded state for long periods. The results for Test 3 depict that although the DAG is in a degraded state, it can still handle the peak messaging load, ensuring that the design can handle one host failure. For Test 4, the results clearly indicate that the read and write latencies for both database and logs are well within the prescribed thresholds. For both Test 3 and Test 4, because MBX2 hosts more databases than MBX1, the resource utilization by MBX2 is greater than that of MBX1.

The host failure considered in the test plan also takes into account unavailability of one of the HUBCAS VMs. Thus, Test 3 and Test 4 were executed with only two HUBCAS VMs to observe the message processing performance of the solution. Table 12 shows the results for HUBCAS1 and HUBCAS2 for both Test 3 and Test 4 respectively.

Table 12. Test Results for VMs Running HUBCAS Server Roles In a Failed Host Configuration

Performance Metrics	Thresholds	Test 3 (Peak load)		Test 4 (Normal load)	
		HUBCAS1	HUBCAS2	HUBCAS1	HUBCAS2
%Processor Utilization	< 70 %	33.149	37.855	16.720	18.561
%Memory utilization (GB)		21.133	21.391	18.690	19.813
RPC Averaged Latency	< 250 ms	5.983	5.755	5.417	5.503
RPC Requests	< 40	5.640	5.106	1.144	1.365
Aggregate Delivery	<3000	2.290	2.081	0.742	1.092
Active Mailbox Delivery Queue Length	<250	1.951	1.951	0.739	1.043
Retry Mailbox Delivery Queue Length	<100	0.090	0.130	0.003	0.049

With a failed host, the performance results with only two HUBCAS VMs indicated that solution design can handle peak messaging load as well as normal messaging load, staying well within the prescribed thresholds. This indicates that from the physical hosts’ perspective, a steady state configuration requires three hosts. However, the design can still handle the peak loads with one host failure, in which case the VMs continue to run on just two physical hosts. Considering the resource utilization for steady state and failed host configurations, the physical hosts running VMs for Exchange server roles can potentially run other workloads, as long as the resource requirements for Exchange server roles are satisfied.

Verifying the VM Placements

This section describes the verification of the DRS rule implementations with Active System 800v, to ensure appropriate placements of VMs running Exchange server roles in the cluster. To verify the ‘Separate VMs’ rule for the mailbox server VM, one of the Mailbox server VMs MBX3 is intentionally migrated to a host already running another mailbox VM - MBX2. Within few minutes, VMware DRS detects the violation of the rule and moves MBX3 back to its original host. Similar verification was performed with a separate VM rule with HUBCAS VMs. VMware DRS detected the violation of the rule and moved the migrated HUBCAS VM back to its original host.

Assessing the Impact of Single DAG Node Failure

This section explains the benefit of Active System 800v in improving Exchange high availability. Whenever there is a DAG node failure, the Exchange Active Manager finds that one of the DAG members has failed, and hence the databases mounted on that DAG member are no longer accessible, so the passive copies of those databases need to be activated. Although the remaining available DAG nodes are sized to take care of DAG node failure, the resource utilization of the functional DAG nodes increases and the DAG enters a degraded state. It is recommended that the DAG does not remain in a degraded state for long periods.

For Exchange Server running in an Active System 800v environment, whenever one of the DAG nodes experiences a failure, the VMware HA and DRS rules help move the failed mailbox VM appropriately to a host, capable of hosting the DAG node VM, helping the failed DAG node to recover and rejoin the Windows cluster and the DAG networks. Once the necessary Exchange services are ready, the recovered DAG node joins the DAG. The logs are replayed on the recovered DAG node to change the state of corresponding database copies to *Healthy*. At this point, the DAG is ready for redistribution of the active databases, to regain the steady state DAG configuration. The following sections describe the verification of such a failure.

Failure Scenario

A steady state DAG configuration and the failed DAG configuration are described in Table 13.

Table 13. Steady State and Failed DAG Configurations

DAG configurations	DAG state	Exchange database distribution
Steady state DAG configuration	Host-1: MBX1, HUBCAS1 Host-2: MBX2, HUBCAS2 Host-3: MBX3, HUBCAS3	MBX1: DB1, DB4, DB7 MBX2: DB2, DB5, DB8 MBX3: DB3, DB6, DB9
Failed host DAG configuration	Host-1: MBX1, HUBCAS1 Host-2: MBX2, HUBCAS2 Host-3: FAILED	MBX1: DB1, DB4, DB7, DB3 MBX2: DB2, DB5, DB8, DB6, DB9
Recovered host DAG configuration	Host-1: MBX1, HUBCAS1 Host-2: MBX2, HUBCAS2 Host-3: FAILED Host-4: MBX3, HUBCAS3	MBX1: DB1, DB4, DB7, DB3 MBX2: DB2, DB5, DB8, DB6, DB9 MBX3: Ready for mounting database copies.

The failure considered in the study is a server hardware failure. The Host-3 running MBX3 and HUBCAS3 VMs in the cluster is brought down using a hard shutdown by holding the power button on the server, simulating a hardware failure. The Exchange Active Manager finds that the DAG node MBX3 has failed, and the databases DB3, DB6, and DB9 mounted on this DAG member are no longer accessible. The Active Manager then activates the passive copy of DB3 on MBX1 and passive copies of DB6 and DB9 on MBX2. Results of the Load Generator tests showed, MBX1 and MBX2 best perform in terms of IOPS and resource utilization when each of them has three databases mounted. After the failure, clearly the DAG is in a degraded state, as MBX1 is handling IOPS for four databases, while MBX2 is handling IOPS for

five. For the given cluster configuration, failure of Host-3 also brings down HUBCAS3. The remaining two HUBCAS server VMs help keep up the email service through the NLB cluster.

The failed MBX3 needs to be recovered in order to redistribute and activate DB3, DB6, and DB9. The HA and the DRS rules implemented in Active System 800v environment help recover MBX3 and HUBCAS3 by moving these VMs to Host-4, which is an available host in the cluster capable of satisfying the virtual resource requirements of MBX3 and HUBCAS3 along with the VM placement policies. Once the necessary Exchange services are ready, MBX3 joins the Windows cluster and DAG networks as the host networking is consistent across all the hosts. The logs are copied to and replayed on MBX3 to change the state of database copies of DB3, DB6, and DB9 hosted by MBX1 and MBX2 to *Healthy*. The databases are ready for re-distribution in order to restore the DAG's original configuration or to account for another host failure. To summarize the scenario: Upon Host-3 failure, Exchange native high availability acts first as the Active Manager detects the failure at the application level, including the cluster nodes. Consequently, VMware HA detects the server failure. VMware HA and DRS identify a host (Host-4) to be suitable for running MBX3 and HUBCAS3 and move them to that available host. It should be noted that as long as the resource requirements of MBX3 and HUBCAS3 are met Host-4 can run other additional workloads too. The failed DAG node rejoins the DAG and Windows cluster and is ready for redistributing the active databases. HUBCAS3 also rejoins the NLB cluster.

Given that the database copies associated with MBX3 reach the *Healthy* state, MBX3 is ready to mount the databases it is assigned to. This means that, in a successive host failure, MBX3 can contribute to mount the databases and continue to provide email service.

Database Redistribution

So far, Active System 800v environment has helped in recovering MBX3, the failed DAG node, without any manual intervention, up to a point at which the concerned administrator just needs to redistribute the active databases. Redistribution of the active databases can also be accomplished by leveraging the `RedistributeActiveDatabases.ps1` script provided with Exchange Server distribution. In the study, during the DAG configuration, the active database copies are configured with an activation preference of 1, and passive copies with an activation preference of 2.

Given that with the recovered DAG state, the databases that failed during the DAG node failure regained their *Healthy* state, the `RedistributeActiveDatabases.ps1` can be used with a `BalanceDbsByActivationPreference` switch and forced confirmation `-Confirm:$false` to redistribute the databases.

NOTE: Administrators should verify that the required host iSCSI and Exchange services on the recovered node have started and the databases are *Healthy* before they are redistributed using `RedistributeActiveDatabases.ps1` script.

Summary

Dell's Active System 800v is designed to provide a highly reliable and scalable platform for virtualization, while simplifying infrastructure management. The incorporation of datacenter standards, like DCB with converged networking, enables Active System 800v to provide performance benefits for enterprise workloads. This guide describes how leveraging Dell's Active System 800v can enable organizations to maximize the value of their infrastructure investments for Exchange Server 2010. Exchange DAGs are a highly recommended method to achieve application-level data protection, including data redundancy and application awareness. Dell's Active System 800v helps increase Exchange uptime while working in co-ordination with Exchange DAG and easing management of Exchange ecosystem while considering the availability requirements for Exchange deployment up to 5,000 mailboxes. Other benefits, such as workload mobility and VM placements with the help of DRS rules, provide an additional layer of security not present in the physical deployments.

This guide explained a design framework to size the various Exchange server roles on Active System 800v. It then explained implementation of a reference configuration consisting of 5,000 users, leveraging the structured approach of the design framework. It should be noted that the same design framework and guidelines from this guide can be leveraged to architect and deploy a solution for custom sized Exchange solution consisting of users up to 5,000.

The guide also provided proof points for performance of Exchange server with Active System 800v and summarized the benefits Exchange can gain from it. It described how Active System 800v works in tandem with the Exchange native HA to help facilitate business continuity for an enterprise email system. Exchange on Active System 800v presents an advantage to Exchange Server by minimizing the management overhead, while improving availability to the application and keeping up with application performance as prescribed by Microsoft.

References

- Dell Email Solutions: <http://www.dell.com/exchange>
- Introduction to Microsoft Exchange Server 2010 sizing: <http://i.dell.com/sites/content/business/solutions/whitepapers/en/Documents/exchange-server-2010-sizing.pdf>
- Mailbox Server Processor Capacity Planning: <http://technet.microsoft.com/en-us/library/ee712771.aspx>
- Understanding Mailbox Database Cache: <http://technet.microsoft.com/en-us/library/ee832793.aspx>
- Understanding Mailbox Database and Log Capacity Factors: <http://technet.microsoft.com/en-us/library/ee832796.aspx>
- Understanding Database and Log Performance Factors: <http://technet.microsoft.com/en-us/library/ee832791>
- Understanding Server Role Ratios and Exchange Performance: <http://technet.microsoft.com/en-us/library/dd346701.aspx>
- Understanding the Impact of Multi-Role Configurations on Scalability: <http://technet.microsoft.com/en-us/library/dd298121.aspx>
- Consulting Services: <http://content.dell.com/us/en/enterprise/by-service-type-it-consulting-communication-collaboration.aspx>
- ProSupport: <http://content.dell.com/us/en/business/d/services/smb-prosupport-for-software-us>
- [Microsoft Lync Server 2010 Implementation on Active System 800v.](#)
- [Microsoft SharePoint Server 2010 Implementation on Active System 800v.](#)

Appendix A: Sizing Virtual CPU for Mailbox Server Role

Microsoft recommends per mailbox database cache, IOPS and CPU estimates based on mailbox user profile and message activity as shown in Table 14. For more details, please refer to *Mailbox Server Processor Capacity Planning* under References.

Table 14. Per Mailbox Database Cache, IOPS and CPU Estimates

Messages sent or received per mailbox per day	Database cache per mailbox in megabytes (MB)	Single database copy (stand-alone) with estimated IOPS per mailbox	Multiple database copies (mailbox resiliency) with estimated IOPS per mailbox	Megacycles for active mailbox or stand-alone mailbox	Megacycles for passive mailbox
50	3	0.06	0.05	1	0.15
100	6	0.12	0.1	2	0.3
150	9	0.18	0.15	3	0.45
200	12	0.24	0.2	4	0.6

To identify the right processor, the Exchange Server 2010 megacycles requirement must be calculated and compared with the megacycles of the processor offered by the server. The total megacycles required are calculated as:

Total Megacycles required = total active mailbox megacycles required + total local passive megacycles required + total remote passive megacycles required

Exchange Server 2010 introduced DAG, a feature that provides application native high availability through replication mechanism. This mechanism gives rise to the additional database copies. Microsoft recommends increasing the megacycles per active mailbox by ten percent for each additional database copy after the one active copy. The example in consideration has two-copy DAG, implying one additional copy. While calculating the total megacycles we consider the worst case scenario where 100% of the active mailbox copies are hosted on a single server.

Total active mailbox megacycles required = [Total number of active mailboxes x Megacycles per mailbox for active mailbox database (shown in Table-14)] x [1 + (number of additional copies x 10 / 100)]

Once the required megacycles for a given profile are identified, based on the type of the physical processor being virtualized and the effective megacycles required, the number of virtual processors can be found. For more details, please refer the mailbox server CPU sizing guidelines.

Microsoft recommends the following guidelines for peak processor utilization for the mailbox server role:

- Standalone servers with only the mailbox server role must be designed to not exceed 70% utilization during peak utilization period.
- For solutions leveraging mailbox resiliency, like DAG, Microsoft recommends that the configuration must not exceed 80% utilization after a single or double member server failure when the server only has the mailbox role installed.
- While deploying multiple roles on the server, the mailbox role should be designed not to exceed 35%. If deploying multiple roles on the server, the mailbox role should be designed not to exceed 40%.

Appendix B: Sizing Virtual CPU for Hub Transport and Client Access Server Roles

Table 15 describes the required processor core ratios. For more details, please refer to ‘*Understanding Server Role Ratios and Exchange Performance*’ under References.

Table 15. Mailbox To (Hub Transport + Client Access) Processor Core Ratio

Number of Client Access and Hub Transport servers	Recommended processor core ratio
Mailbox server role: Client Access and Hub Transport combined server roles	1:1
Mailbox server role: Client Access	4:3
Mailbox server role: Hub Transport	7:1 (without antivirus scanning on Hub Transport) 5:1 (with antivirus scanning on Hub Transport)

Based on these ratios, the virtual CPU requirements for the Hub Transport and Client Access Server roles can be identified.

Appendix C: Sizing Virtual CPU for Other Exchange Server Roles

VMware recommends the following for minimum and maximum processor allocations to various Exchange server roles:

Table 16. Processor Configurations for Various Exchange Server Roles

Exchange 2010 Server roles	Minimum processor cores	Recommended maximum processor cores
Mailbox	2	12
Hub Transport	1	12
Client Access	1	12
Edge Transport	1	12
Unified Messaging	1	12
Client Access and Hub Transport Collocated	2	12
Multi-role (Mailbox, Client Access and Hub transport running on same server)	2	24

Appendix D: Sizing Memory for Mailbox Server Role

Based on the mailbox message profile, the database cache can be determined from Table 17, per Microsoft guidelines. For more details, please refer to ‘*Understanding Mailbox Database Cache*’ under References.

Table 17. Database Cache Requirements

Messages sent or received per mailbox per day (about 75 KB average message size)	Database cache per user (MB)
50	3 MB
100	6 MB
150	9 MB
200	12 MB

Once the database cache per user is identified, the total database cache required per server is calculated as follows:

Total database cache required per server = Number of active mailboxes per server x average database cache per Mailbox

The next step is to determine the memory of the entire mailbox server. Based on the lookup table provided below and the nature of the mailbox server role deployment (collocated or standalone), the required memory can be easily identified.

Table 18. Mailbox Memory Recommendations by Microsoft Based on Database Cache

Server physical memory (RAM)	Database cache size (Mailbox role only)	Database cache size: Multiple-role (for example, Mailbox + Hub Transport)
2GB	512 MB	Not supported
4GB	1 GB	Not supported
8GB	3.6 GB	2 GB
16GB	10.4 GB	8 GB
24GB	17.6 GB	14 GB
32GB	24.4 GB	20 GB
48GB	39.2 GB	32 GB
64GB	53.6 GB	44 GB
96GB	82.4 GB	68 GB

Apart from the server roles, the hypervisor and guest memory requirements should be taken in to consideration.

Summary of Memory Requirements for Various Exchange Server Roles

Memory recommendations from Microsoft pertaining to the different Exchange server roles are summarized in Table 19.

Table 19. Memory Recommendations

Exchange 2010 Server Role	Minimum Supported	Recommended
Edge Transport	4GB	1GB per core (4GB minimum)
Hub Transport	4GB	1GB per core (4GB minimum)
Client Access	4GB	2GB per core (8GB minimum)
Unified Messaging	4GB	2GB per core (4GB minimum)
Mailbox	4GB	4GB plus 3-30MB/mailbox: This variable is based on the user profile.
Client Access/Hub Transport combined role (Client Access and Hub Transport server roles running on the same physical server)	4GB	2GB per core (8GB minimum)
Multiple roles (combinations of Hub Transport, Client Access, and Mailbox server roles)	10GB	10GB plus 3-30MB per mailbox (4-core server) 14GB plus 3-30MB per mailbox (8-core server) 18GB plus 3-30MB per mailbox (12-core server) 22GB plus 3-30MB per mailbox (16 core server) 30GB plus 3-30MB per mailbox (24-core server) This variable is based on the user profile.