



Microsoft Exchange 2010 on VMware Best Practices Guide

© 2011 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at <http://www.vmware.com/download/patents.html>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com

Contents

1.	Introduction	7
1.1	Purpose.....	7
1.2	Target Audience.....	7
1.3	Scope	8
2.	VMware ESX Host Best Practices for Exchange.....	9
2.1	CPU Configuration Guidelines	9
2.2	Memory Configuration Guidelines.....	10
2.3	Storage Virtualization	12
2.4	Networking Configuration Guidelines	16
3.	Exchange Performance on vSphere.....	19
3.1	Overview	19
3.2	Key Performance Considerations	20
3.3	Performance Testing.....	20
3.4	Ongoing Performance Monitoring and Tuning	21
3.5	vSphere Performance Features and Exchange.....	22
4.	Exchange 2010 Capacity Planning	23
4.1	Capacity Planning Process Overview	23
4.2	Mailbox Server: Planning for CPU	24
4.3	Mailbox Server: Planning for Memory.....	27
4.4	Mailbox Server: Planning for Storage	29
4.5	Hub Transport and Client Access Server Planning	30
4.6	Scaling Exchange for the Enterprise.....	32
4.7	vSphere Limitations	34
5.	Sizing Examples	35
5.1	Standalone Mailbox Server – 16,000 Users (150 sent/received)	35
5.2	Clustered Mailbox Server – 16,000 Users (150 sent/received)	50
6.	vSphere Enhancements for Deployment and Operations.....	66
6.1	VMware vMotion, DRS, and VMware HA	66
6.2	Templates	67
6.3	VMware vCenter Operations.....	67
6.4	VMware vCenter Site Recovery Manager	68
6.5	VMware vCenter AppSpeed	70

List of Tables

Table 1. VMFS and Raw Disk Mapping Trade-offs	14
Table 2. Performance Testing Summary	21
Table 3. Performance Counters of Interest to Exchange Administrators	22
Table 4. Exchange Server Minimums and Recommended Maximums	25
Table 5. Per Mailbox Database Cache, IOPS, and CPU Estimates Based on User Profile and Message Activity*	26
Table 6. Memory Configurations for Exchange 2010 Servers Based on Installed Server Roles*	27
Table 7. Per Mailbox Database Cache, IOPS, and CPU Estimates Based on User Profile and Message Activity*	28
Table 8. Determining Total Memory	29
Table 9. Recommended Server Role Ratios Based on Processor Core*	30
Table 10. Recommended Server Role Memory Based on Processor Core*	31
Table 11. Building block CPU and RAM requirements for mailboxes with 150 messages sent/received per day	32
Table 12. Mailbox Server CPU Recommendations*	35
Table 13. Mailbox Server Database Cache Requirements*	37
Table 14. Default Mailbox Database Cache Sizes*	37
Table 15. Environment Configuration	38
Table 16. Data Configuration	38
Table 17. Mailbox Configuration	39
Table 18. Backup Configuration	39
Table 19. Storage Configuration	40
Table 20. Processor Configuration	40
Table 21. Mailbox Server Storage Results	40
Table 22. 4,000-User Building Block Requirements (150 sent/received)	41
Table 23. Exchange Virtual Machine Configuration	42
Table 24. Recommended Server Role Ratios Based on Processor Core*	44
Table 25. Recommended Server Role Ratios Based on Processor Core*	44
Table 26. Example Exchange Server Role Resource Requirements	46
Table 27. Example Exchange Virtual Machine Configuration	47
Table 28. Example Exchange Virtual Machine Distribution	48
Table 29. Example ESXi Host Hardware Configuration Table	49
Table 30. Mailbox Server CPU Recommendations*	50
Table 31. Mailbox Server Database Cache Requirements*	52
Table 32. Default mailbox database cache sizes*	53
Table 33. Environment Configuration	54

Table 34. Data Configuration	54
Table 35. Database Copy Configuration	55
Table 36. Database Configuration	55
Table 37. Mailbox Configuration	55
Table 38. Backup Configuration.....	56
Table 39. Storage Configuration	56
Table 40. Processor Configuration	56
Table 41. Mailbox Server Storage Results	57
Table 42. 16,000 Active Mailboxes (150 sent/received) DAG Node Requirements	58
Table 43. Exchange Virtual Machine Configuration.....	59
Table 44. Recommended Server Role Ratios Based on Processor Core*	60
Table 45. Recommended Server Role Ratios Based on Processor Core*	60
Table 46. Example Exchange Server Role Resource Requirements	62
Table 47. Exchange Virtual Machine Configuration.....	63
Table 48. Exchange Virtual Machine Distribution	64
Table 49. Example ESXi Host Hardware Configuration Table	65

List of Figures

Figure 1. Virtual Machine Memory Settings	11
Figure 2. VMware Storage Virtualization Stack	13
Figure 3. Storage Multipathing Requirements for vSphere.....	13
Figure 4. One Versus Many Virtual Machines in a LUN	16
Figure 5. Sample Network Layout for Exchange Environment	18
Figure 6. Passive Database Overhead	24
Figure 7. Building Block Virtual Machine Interaction with Shared Storage.....	43
Figure 8. Initial Virtual Machine Placement.....	49
Figure 9. Mailbox Server Virtual Machine Interaction with Shared Storage	60
Figure 10. Initial Virtual Machine Placement.....	65
Figure 11. VMware vCenter Operations	67
Figure 12. VMware vCenter Site Recovery Manager	68
Figure 13. VMware vCenter AppSpeed	70

1. Introduction

Email has become one of the most critical applications in an organization's IT infrastructure. Organizations increasingly rely on messaging tools for individual and organizational effectiveness. As a result, messaging administrators face a constant challenge as they continually seek to manage the conflicting demands of availability, agility, and cost.

Microsoft Exchange is the most widely used email system in the world. Its operational and performance characteristics are well understood, and best practices for design, deployment, and operations are readily accessible. Exchange continues to evolve through enhanced features and functionality, and through previous limitations addressed with each successive new version.

With its release of Exchange Server 2010, Microsoft added many features that improve messaging performance, reliability, and scalability. These provide a major step forward. However, Exchange Server 2010 is still subject to many of the shortcomings inherent in most applications running directly on physical hardware, such as hardware platform dependence, under-utilization of server computing resources, lack of flexibility to respond to changing workloads, and heavy costs associated with maintaining disaster recovery, test, and development environments. The architectural improvements in Exchange Server 2010 cannot fully address these limitations.

The ideal platform for Exchange would adapt easily to changing workloads, provide flexibility to accommodate changing demands on an organization's IT infrastructure, remain reliable and resilient despite system outages, and improve both staff and infrastructure hardware effectiveness. A new operational platform based on VMware vSphere® can accomplish these goals.

1.1 Purpose

This guide provides best practice guidelines for deploying Exchange Server 2010 on vSphere. The recommendations in this guide are not specific to any particular set of hardware or to the size and scope of any particular Exchange implementation. The examples and considerations in this document provide guidance, but do not represent strict design requirements as the flexibility of Exchange Server 2010 on vSphere allows for a wide variety of valid configurations.

1.2 Target Audience

This guide assumes a basic knowledge and understanding of vSphere and Exchange Server 2010.

- Architectural staff can use this document to gain an understanding of how the system will work as a whole as they design and implement various components.
- Engineers and administrators can use this document as a catalog of technical capabilities.
- Messaging staff can use this document to gain an understanding of how Exchange might fit into a virtual infrastructure.
- Management staff and process owners can use this document to help model business processes to take advantage of the savings and operational efficiencies achieved with virtualization.

1.3 Scope

The scope of this document is limited to the following topics:

- VMware ESX Host Best Practices for Exchange – This section provides best practice guidelines for preparing the vSphere platform for running Exchange Server 2010. Guidance is included for CPU, memory, storage, and networking.
- Exchange Performance on vSphere – This section provides background information on Exchange Server 2010 performance in a virtual machine. It also provides information on official VMware partner testing and guidelines for conducting and measuring internal performance tests.
- Exchange 2010 Capacity Planning – Sizing Exchange 2010 to run in a virtual machine follows many of the same best practices as sizing on physical servers; however, with the introduction of new Exchange 2010 features (Database Availability Groups), the Capacity Planning process has changed significantly. This section walks through this new process.
- Sizing Examples – In this section, we apply the new Capacity Planning process to two sample configurations, one with Database Availability Groups and one without.
- vSphere Enhancements for Deployment and Operations – This section provides a brief look at vSphere features and add-ons that enhance deployment and management of Exchange 2010.

The following topics are out of scope for this document, but may be addressed in other documentation in this solution kit:

- Design and Sizing Examples – This information can be found in the Microsoft Exchange 2010 on VMware: Design and Sizing Examples document included in this Solution Kit, which expands upon the examples in this guide by showing how the Capacity Planning process works for small, medium, and enterprise configurations.
- Availability and Recovery Options – Although this document briefly covers VMware features that can enhance availability and recovery, a more in-depth discussion of this subject is covered in *Microsoft Exchange 2010 on VMware: Availability and Recovery Options* (included in this solution kit).

It is important to note that this and other guides in this Solution Kit are limited in focus to deploying Exchange on vSphere. Exchange deployments cover a wide subject area, and Exchange-specific design principles should *always* follow Microsoft guidelines for best results.

2. VMware ESX Host Best Practices for Exchange

A solidly designed ESXi host platform is crucial to the successful implementation of enterprise applications such as Exchange. Before we address best practices specific to the Exchange application, the following sections outline general best practices for designing vSphere.

2.1 CPU Configuration Guidelines

2.1.1 Physical and Virtual CPUs

VMware uses the terms virtual CPU (vCPU) and physical CPU to distinguish between the processors within the virtual machine and the underlying physical x86/x64-based processor cores. Virtual machines with more than one virtual CPU are also called SMP (symmetric multi-processing) virtual machines. The virtual machine monitor (VMM) is responsible for virtualizing the CPUs. When a virtual machine starts running, control transfers to the VMM, which is responsible for virtualizing guest OS instructions.

2.1.2 Virtual SMP

VMware Virtual Symmetric Multi-Processing (Virtual SMP) enhances virtual machine performance by enabling a single virtual machine to use multiple physical processor cores simultaneously. vSphere supports use of up to 32 virtual CPUs per virtual machine. The biggest advantage of an SMP system is the ability to use multiple processors to execute multiple tasks concurrently, thereby increasing throughput (for example, the number of transactions per second). Only workloads that support parallelization (including multiple processes or multiple threads that can run in parallel) can really benefit from SMP.

The virtual processors from SMP-enabled virtual machines are co-scheduled. That is, if physical processor cores are available, the virtual processors are mapped one-to-one onto physical processors and are then run simultaneously. In other words, if one vCPU in the virtual machine is running, a second vCPU is co-scheduled so that they execute nearly synchronously. Consider the following points when using multiple vCPUs:

- Simplistically, if multiple, idle physical CPUs are not available when the virtual machine wants to run, the virtual machine will remain in a special wait state. The time a virtual machine spends in this wait state is called *ready time*.
- Even idle processors perform a limited amount of work in an operating system. In addition to this minimal amount, the ESXi host manages these “idle” processors, resulting in some additional work by the hypervisor. These low-utilization vCPUs compete with other vCPUs for system resources.

VMware ESX™ 4 provided several changes to the CPU scheduler which are detailed in the paper *VMware vSphere 4: The CPU Scheduler in VMware ESX 4* (http://www.vmware.com/files/pdf/perf-vsphere-cpu_scheduler.pdf). For example, in ESX 4, the relaxed co-scheduling algorithm has been refined so that scheduling constraints due to co-scheduling requirements are further reduced. These improvements carried over into vSphere 5 and result in better linear scalability and performance of Exchange workloads, as described in Section 3, Exchange Performance on vSphere, while reducing inefficiencies introduced by idle vSMP virtual machines. Consequently, in vSphere, the larger 8-way and 12-way virtual machines exhibit great scalability and are a much more viable option if there is a requirement to scale up versus scale out.

Consequently, VMware recommends the following practices:

- Only allocate multiple vCPUs to a virtual machine if the anticipated Exchange workload can truly take advantage of all the vCPUs.
- If the exact workload is not known, size the virtual machine with a smaller number of vCPUs initially and increase the number later if necessary.
- For performance-critical Exchange virtual machines (production systems), the total number of vCPUs assigned to all the virtual machines should be equal to or less than the total number of cores on the ESXi host machine.

While larger virtual machines are possible in vSphere, VMware recommends reducing the number of virtual CPUs if monitoring of the actual workload shows that the Exchange application is not benefitting from the increased virtual CPUs. For more background information, see the “ESXi CPU Considerations” section in the white paper *Performance Best Practices for VMware vSphere 5* (http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf).

Setting a CPU Reservation sets a guaranteed CPU allocation for the virtual machine. This practice is generally not recommended because the reserved resources are not available to other virtual machines and flexibility is often required to manage changing workloads. However, SLAs and multi-tenancy may require a guaranteed amount of compute resource to be available. In these cases, reservations ensure these requirements are met. VMware has conducted tests on virtual CPU over-commitment with SAP and SQL, showing that the performance degradation inside the virtual machines is linearly reciprocal to the over-commitment. As the performance degradation is “graceful,” any virtual CPU over-commitment can be effectively managed by using DRS and vMotion to move virtual machines to other ESXi hosts to obtain more processing power.

2.1.3 Hyper-Threading

Hyper-threading technology (recent versions of which are called symmetric multithreading, or SMT) allows a single physical processor core to behave like two logical processors, essentially allowing two independent threads to run simultaneously. Unlike having twice as many processor cores—which can roughly double performance—hyper-threading can provide anywhere from a slight to a significant increase in system performance by keeping the processor pipeline busier. For example, an ESXi host system enabled for SMT on an 8-core server will see 16 threads that appear as 16 logical processors.

2.2 Memory Configuration Guidelines

This section provides guidelines for allocation of memory to Exchange virtual machines. The guidelines outlined here take into account vSphere memory overhead and the virtual machine memory settings.

2.2.1 ESXi Memory Management Concepts

vSphere virtualizes guest physical memory by adding an extra level of address translation. Shadow page tables make it possible to provide this additional translation with little or no overhead. Managing memory in the hypervisor enables the following:

- Memory sharing across virtual machines that have similar data (same guest operating systems).
- Memory over-commitment, which means allocating more memory to virtual machines than is physically available on the ESXi host.
- A memory balloon technique whereby virtual machines that do not need all the memory they have been allocated give memory to virtual machines that require additional allocated memory.

For more details about vSphere memory management concepts, consult the *VMware vSphere Resource Management Guide* (http://www.vmware.com/pdf/vsphere4/r40_u1/vsp_40_u1_resource_mgmt.pdf).

2.2.2 Virtual Machine Memory Concepts

Figure 1 illustrates the use of memory settings parameters in the virtual machine.




Figure 1. Virtual Machine Memory Settings

The vSphere memory settings for a virtual machine include the following parameters:

- Configured memory = memory size of virtual machine assigned at creation.
- Touched memory = memory actually used by the virtual machine. vSphere only allocates guest operating system memory on demand.
- Swappable = virtual machine memory that can be reclaimed by the balloon driver or by vSphere swapping. Ballooning occurs before vSphere swapping. If this memory is in use by the virtual machine (touched and in use), the balloon driver will cause the guest operating system to swap. Also, this value is the size of the per-virtual machine swap file that is created on the VMware Virtual Machine File System (VMFS) file system (VSWP file).
- If the *balloon driver* is unable to reclaim memory quickly enough, or is disabled or not installed, vSphere forcibly reclaims memory from the virtual machine using the *VMkernel swap file*.

2.2.3 Allocating Memory to Exchange Virtual Machines

Microsoft has developed a thorough sizing methodology for Exchange server that has matured over the last couple of versions of Exchange. VMware recommends using the memory sizing guidelines set by Microsoft. This methodology is discussed in detail in Section 4.2, Mailbox Server: Planning for CPU. Simplistically the amount of memory required for an Exchange server is driven by its role and, if it is a mailbox server, the number of mailboxes on that server. From the perspective of VMware specifics the architect should consider the VMM memory requirements on top of the memory requirements for the Exchange server itself.

As Exchange servers are memory intensive and performance is often a key factor (for example, in production environments), VMware recommends the following practices:

- Do not over-commit memory on ESXi hosts running Exchange workloads. For production systems, it is possible to enforce this policy by setting the memory reservation to the configured size of the virtual machine. Also note that:
 - Setting memory reservations may limit vMotion. A virtual machine can only be migrated if the target ESXi host has free physical memory equal to or greater than the size of the reservation.
 - Setting the memory reservation to the configured size of the virtual machine results in a per-virtual machine vmkernel swap file of zero bytes that will consume less storage and help increase performance by eliminating ESXi host-level swapping. The guest operating system within the virtual machine maintains its own separate swap/page file.
 - Reservations are generally only recommended when it's possible memory may become over-committed on hosts running Exchange virtual machines, when SLAs dictate that memory be "guaranteed," or when there is a desire to reclaim space used by a virtual machine swap file.

- It is important to “right-size” the configured memory of a virtual machine. Memory is wasted if the Exchange virtual machines are not utilizing the configured memory. ESXi performance counters can be used to determine actual memory usage (see Section 3, Exchange Performance on vSphere).
- Do not disable the balloon driver (which is installed with VMware Tools).
- Enable DRS to balance workloads in the ESXi cluster. DRS and reservations can guarantee critical workloads have the resources they require to operate optimally.
- To minimize guest operating system (OS) swapping, the configured memory size of the virtual machine should be greater than the average memory usage of Exchange running in the guest OS. If the Exchange virtual machine needs more memory than has been allocated, the guest OS paging/swapping mechanisms are invoked as in normal, native operations. Memory and swap/page file configuration for Exchange virtual machines follow the same guidelines as for native environments. In general, these should be set to minimize any guest OS swapping.

2.2.4 Advanced Memory Management

The guidelines described above are purposely conservative to avoid kernel swapping between ESXi and the guest OS—important due to the mission-critical nature of Exchange, which must often meet stringent SLAs, and the memory intensive nature of the application. This best practice can also apply to non-production systems with high performance SLAs for developers and testers who support production environments.

However, it is feasible that once the Exchange workload is known and predictable, if VMware vCenter reports that steady state active memory usage is below the amount of memory on the ESXi host, then the reservation settings may be relaxed to the steady state active memory value. This scenario is discussed in the VMworld® 2009 presentation, TA2627 – *Understand “Host” and “Guest” Memory Usage and Related Memory Management Concepts* (<http://www.vmworld2009.com/docs/DOC-3817>). VMworld account is required for access.

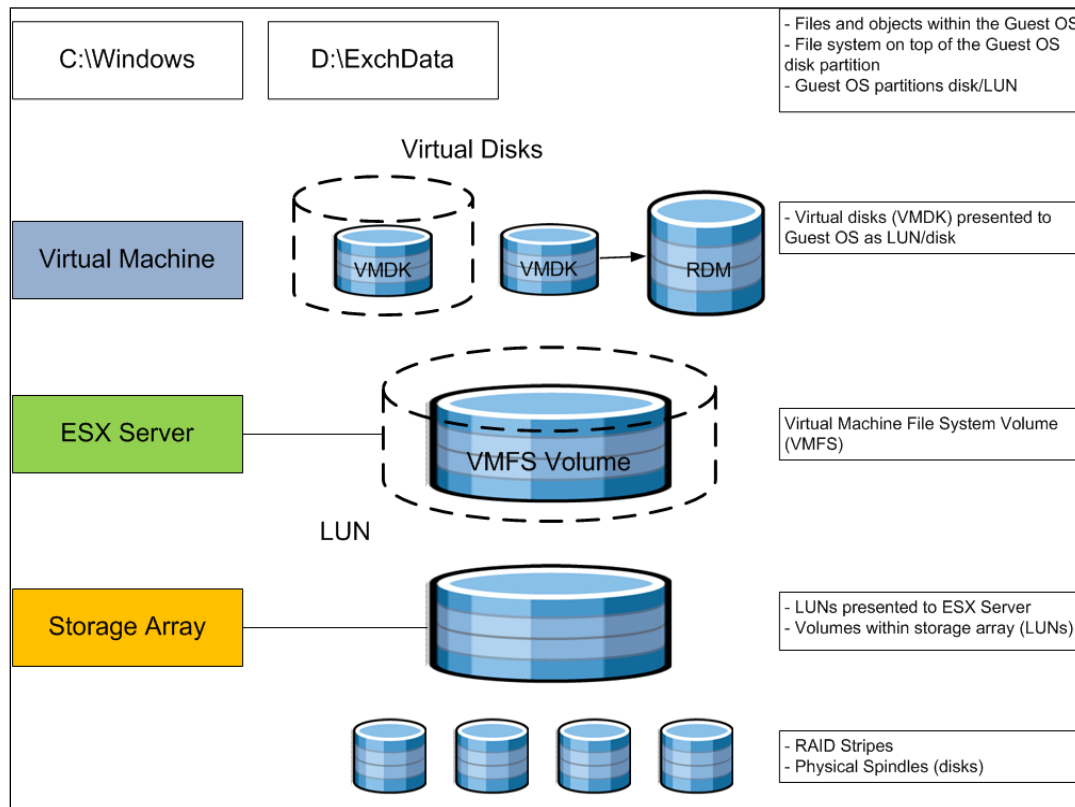
2.3 Storage Virtualization

VMFS is a cluster file system that provides storage virtualization optimized for virtual machines. Each virtual machine is encapsulated in a small set of files and VMFS is the default storage system for these files on physical SCSI disks and partitions. VMware supports Fibre-Channel, iSCSI, and NAS shared-storage protocols.

It is preferable to deploy virtual machine files on shared storage to take advantage of vMotion, VMware High Availability (HA), and VMware Distributed Resource Scheduler (DRS). This is considered a best practice for mission-critical Exchange deployments, which are often installed on third-party, shared-storage management solutions.

VMware storage virtualization can be categorized into three layers of storage technology, as illustrated in Figure 2. The storage array is the bottom layer, consisting of physical disks presented as logical disks (storage array volumes or LUNs) to the layer above, with the virtual environment occupied by vSphere. Storage array LUNs that are formatted as VMFS volumes in which virtual disks can be created. Virtual machines consist of virtual disks that are presented to the guest operating system as disks that can be partitioned and used in file systems.


Figure 2. VMware Storage Virtualization Stack



2.3.1 Storage Multipathing

VMware recommends you set up a minimum of four paths from an ESXi host to a storage array, so the host requires at least two HBA ports.

Figure 3. Storage Multipathing Requirements for vSphere



The terms used in Figure 3 are:

- HBA (Host Bus Adapter) – A device that connects one or more peripheral units to a computer and manages data storage and I/O processing.
- FC (Fibre Channel) – A gigabit-speed networking technology used to build storage area networks (SANs) and to transmit data.
- SP (Storage Processor) – A SAN component that processes HBA requests routed through an FC switch and handles the RAID/volume functionality of the disk array.

2.3.2 Raw Device Mapping

VMFS also supports Raw Device Mapping (RDM). RDM allows a virtual machine to directly access a volume on the physical storage subsystem, and can only be used with Fibre Channel or iSCSI. RDM can be thought of as providing a symbolic link from a VMFS volume to a raw volume. The mapping makes volumes appear as files in a VMFS volume. The mapping file, not the raw volume, is referenced in the virtual machine configuration.

There are no concrete recommendations for using VMFS or RDM in Exchange deployments, although the following table summarizes some of the options and trade-offs. For a more complete discussion, see the *VMware SAN System Design and Deployment Guide* (<http://www.vmware.com/resources/techresources/772>).

Table 1. VMFS and Raw Disk Mapping Trade-offs

VMFS	RDM
<ul style="list-style-type: none"> • Volume can host many virtual machines (or can be dedicated to one virtual machine). • Increases storage utilization, provides better flexibility, easier administration, and management. • Large third-party ecosystem with V2P products to aid in certain support situations. • Does not support quorum disks required for third-party clustering software. This does not apply to configurations where shared storage is not used. For example, Exchange with DAG or CCR. • Fully supports VMware vCenter Site Recovery Manager. • Supports vMotion, HA and DRS. • Virtual disks limited to 2TB. 	<ul style="list-style-type: none"> • Ideal if disks must be dedicated to a single virtual machine. • More LUNs are required, so it is easier to reach the LUN limit of 256 that can be presented to an ESXi host. • Required to leverage array-level backup and replication tools (VSS) integrated with Exchange databases. • RDM volumes can help facilitate migrating Exchange data from physical to new virtual machines by using the LUN swing method. • Required for in-guest clustering (for example, MSCS) between ESXi hosts. Cluster data and quorum disks must be configured with RDM. Configurations not requiring shared storage, for example, Exchange DAG/CCR or MNS (Majority Node Set) are not subject to this requirement. • Fully supports VMware vCenter Site Recovery Manager. • Supports vMotion, HA and DRS. • Supports LUNs up to 64TB.

It is also possible and even advantageous in some circumstances to mix VMFS and RDM in Exchange environments under the following conditions:

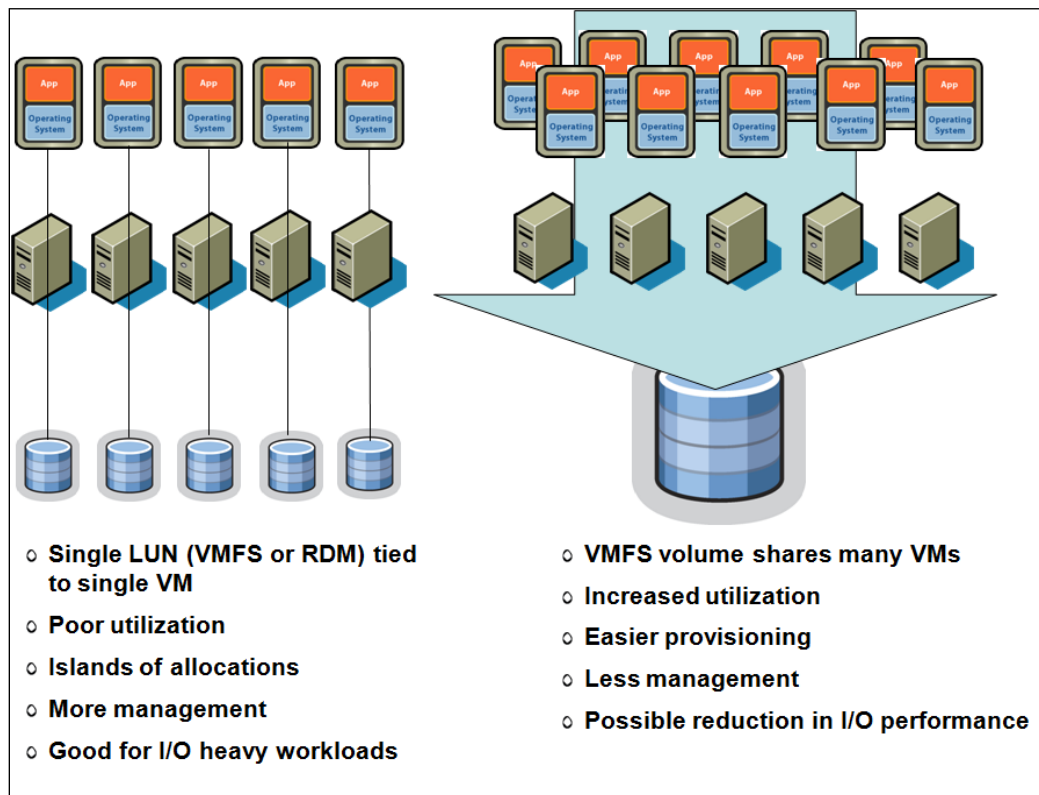
- Where existing systems already make use of third-party storage management software, RDM can be used to leverage practices based on these products such as storage-based backups to disk.
- RDM is required when using third-party clustering software. Configurations not requiring shared storage. For example, Exchange DAG/CCR or MNS (Majority Node Set) are not subject to this requirement.
- RDM is useful for enabling the database portability feature of the Exchange database. Running the database on an RDM volume gives an administrator the option of pointing both virtual machines and physical servers to the same storage. This can be particularly useful in support situations that require problems be reproduced on a physical server.
- Deploying multiple, non-production Exchange systems on VMFS facilitates easier management and administration of template cloning, snapshots, and storage consolidation.
- A mixed storage configuration is viable for an Exchange virtual machine. The guest OS is installed on VMFS and the Exchange database and log files on RDM. VMware template cloning can be used for the guest OS and database files can be managed by third-party storage management software.
- Database datafiles should be spread out over multiple LUNs, similar to those in native setups, following the storage vendor or ISV guidelines for database layout, LUN and spindle configuration.
- Maintain a 1:1 mapping between the number of virtual machines and LUNs to avoid any disk I/O contention.
- A minimum of two HBA adaptors should be configured per ESXi host.
- Follow the guidelines in the “Hardware Storage Considerations” and “Guest Operating Systems” sections of *Performance Best Practices for VMware vSphere 5*.

There are several different shared-storage options available to ESXi (iSCSI, Fibre Channel, NAS, others.); however, Microsoft does not currently support NFS for the Mailbox Server role (clustered or standalone). For Mailbox servers that belong to a Database Availability Group, non-shared storage can be accessed via guest OS-based Software iSCSI Initiators or ESXi host-based virtual disks and RDMs provided that none of the MSCS configurations used require shared storage. Standalone mailbox servers can use iSCSI storage via guest OS-based software iSCSI Initiators or ESXi host-based iSCSI Initiators. To see the latest compatibility list see the *VMware Compatibility Guide* (<http://www.vmware.com/resources/compatibility/search.php>).

2.3.3 Number of Virtual Machines per LUN

The number of virtual machines allocated to a VMFS LUN influences the final architecture. Figure 4 illustrates the concepts and highlights the differences between a one-to-one and many-to-one virtual machine to LUN assignment.

Figure 4. One Versus Many Virtual Machines in a LUN



2.4 Networking Configuration Guidelines

This section covers design guidelines for the virtual networking environment and provides configuration examples at the ESXi host level for Exchange Server 2010 installations.

Note The examples do not reflect design requirements and do not cover all possible Exchange network design scenarios.

2.4.1 Virtual Networking Concepts

The virtual networking layer consists of the virtual network devices through which virtual machines and the service console interface with the rest of the network and users. In addition, ESXi hosts use the virtual networking layer to communicate with iSCSI SANs and NAS storage.

The virtual networking layer includes virtual network adapters and the virtual switches. Virtual switches are the key networking components in vSphere. They are “built to order” at run time and are implemented in much the same way as a modern Ethernet switch, supporting functions equivalent to VLANs based on the IEEE 802.1Q protocol.

2.4.1.1. Virtual Switches and Port Groups

Virtual switches work like Ethernet switches and support VLAN segmentation at the port level. VLANs in vSphere allow logical groupings of switch ports to communicate as if all ports were on the same physical LAN segment. VLANs require tagging of Ethernet frames with the 802.1Q tag (based on IEEE protocol standards), and vSphere enables port-based VLAN tagging based on the switch ports. *VMware ESX Server 3: 802.1Q VLAN Solutions* (http://www.vmware.com/pdf/esx3_vlan_wp.pdf) discusses three different configuration modes to tag:

- Virtual Switch Tagging (VST mode) – Virtual switch port group adds and removes tags.
- Virtual Machine Guest Tagging (VGT mode) – An 802.1Q VLAN trunking driver is installed in the virtual machine.
- External Switch Tagging (EST mode) – External switches perform VLAN tagging so that Ethernet frames moving in and out of the ESXi host are not tagged with VLAN IDs.

The most common configuration is VST mode. VST mode requires provisioning one port group on a virtual switch for each VLAN and attaching the virtual machine's virtual adapter to the port group of the virtual switch. The virtual switch port group tags all outbound frames and removes tags for all inbound frames. It also ensures that frames on one VLAN are isolated from other VLANs. VST mode requires that the physical switch provide a trunk (*trunking* is the technology that allows information from multiple VLANs to be carried over a single link).

Port groups in vSphere are templates for creating virtual ports with a particular set of specifications. In vSphere, there are three types of port group/virtual switch connections:

- Service console port group – vSphere management interface.
- VMkernel port group – vMotion, iSCSI, and/or NFS/NAS networks.
- Virtual machine port group – Virtual machine networks.

More than one connection type can exist on a single virtual switch, or each connection type can exist on its own virtual switch.

2.4.1.2. NIC Teaming

vSphere allows a single virtual switch to be connected to multiple, physical Ethernet adapters using a feature called *NIC teaming*. This provides redundancy and may provide aggregation depending on the upstream switch configuration.

2.4.2 Virtual Networking Best Practices

The standard VMware networking best practices apply to running Exchange on vSphere:

- Allocate separate network adapters/networks for vMotion, VMware FT logging traffic, and ESXi console access management.
- Allocate at least two network adapters for Exchange production traffic to leverage VMware NIC teaming capabilities. Generally, at least four network adapters are recommended per ESXi host.
- Use the VMXNET3 network adapter – This is a paravirtualized device that works only if VMware Tools is installed on the guest operating system. The VMXNET3 adapter is optimized for virtual environments and designed to provide high performance.
- To support VLANs in vSphere, the virtual or physical network must tag the Ethernet frames with 802.1Q tags using virtual switch tagging (VST), virtual machine guest tagging (VGT), or external switch tagging (EST). VST mode is the most common configuration.
- Follow the networking design guidelines in VMworld 2010 session TA8595 - *Virtual Networking Concepts and Best Practices* – this includes designs to efficiently manage multiple networks and redundancy of network adaptors on ESXi hosts.

- Follow the guidelines in the “Hardware Networking Considerations” and “Guest Operating Systems” sections of *Performance Best Practices for VMware vSphere 5* (http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf).

2.4.3 Sample Exchange Virtual Network Configuration

Figure 5 is an example of how a network layout for an Exchange production environment might look.

Figure 5. Sample Network Layout for Exchange Environment

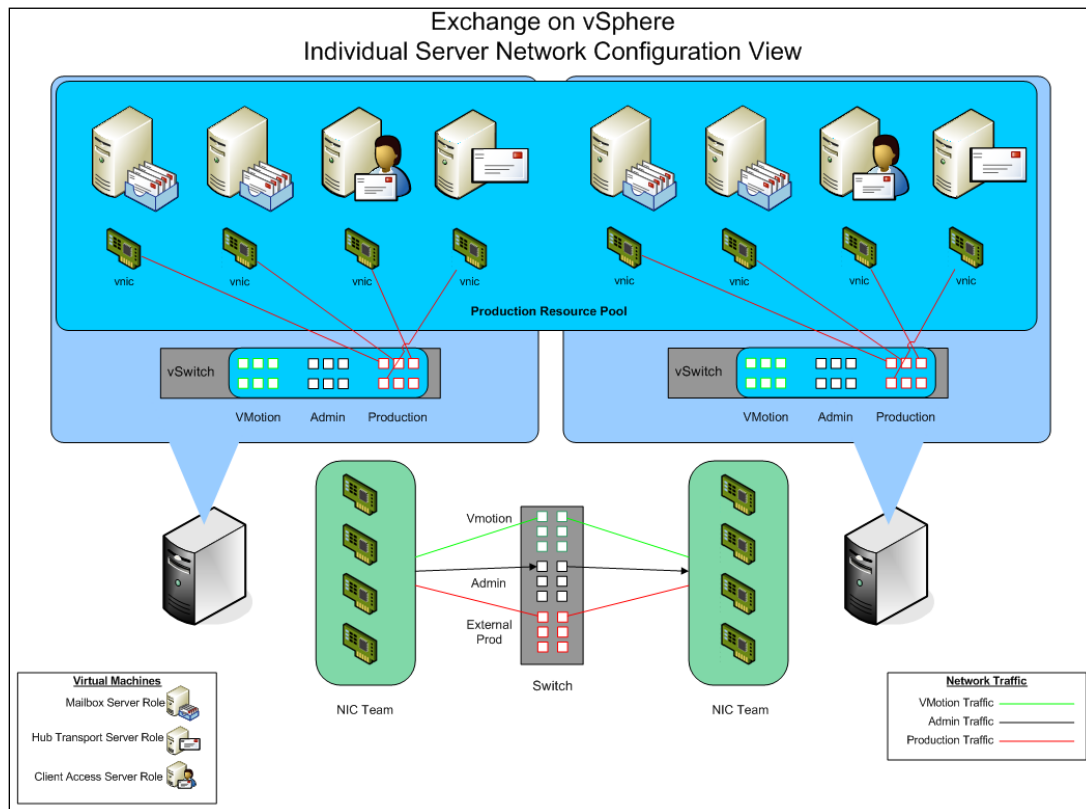


Figure 5 illustrates how networking is handled at the ESXi level. Each ESXi host must have virtual switches architected to handle the type of network traffic that will be assigned to each of the different virtual machines. The figure represents a sample configuration where the production resource pool is split between two physical servers (to reflect redundancy for HA considerations). From a networking perspective, make sure that production environment network traffic remains separate from vMotion and Admin traffic. An effective way to handle this is by introducing VLAN technology to logically separate the traffic.

Each virtual machine acts independently, and remains isolated until networking is configured. What makes the environment different than that in the physical world is that it must have an internal network configured to establish communication between virtual machines residing on the same physical ESXi host. This network traffic is handled through the virtual switch.

Each physical NIC can be configured to connect directly to an assigned VLAN, but the vMotion and Admin networks are not used as heavily as production networks. One practice is to team all the NICs on the ESXi host, connect them to trunk ports on the switch, and use VLAN tagging to direct the traffic at the switch level. This allows for better bandwidth utilization and frees up server capacity for production traffic when the vMotion and Admin VLANs are not in heavy use.

3. Exchange Performance on vSphere

3.1 Overview

Since 2006, VMware and its partners have used testing to successfully demonstrate the viability of running Exchange on the VMware Infrastructure platform. This testing has been confirmed by organizations that have deployed Exchange 2003, 2007, and 2010 in virtualized production environments and now benefit from the considerable operational advantages and cost savings. Some customers have virtualized their entire Exchange 2010 environment and have carefully designed their vSphere infrastructure to accommodate application performance, scalability, and availability requirements.

Exchange Server 2010 is proving to be an even better candidate for virtualization than its predecessors. With up to 50% reduction in disk I/O from Exchange 2007 and many of the mailbox services being offloaded onto the Client Access Server role, the Exchange 2010 Mailbox Server role is lightweight, fast, and ready for virtualization. The shift towards running Exchange in virtual machines is a result of advancements in three key areas:

- Architectural improvements of Exchange Server 2010 have drastically increased memory utilization and reduced disk I/O load by 50 percent or more in many cases, addressing many of the shortcomings found in Exchange 2003 and 2007.
- Advances in server hardware such as multi-core processors, higher memory density, and advances in storage technology are far outpacing the performance requirements for today's applications, including Exchange. Virtualization becomes an effective way to leverage the full power of these systems.
- The advances in Exchange Server 2010 and server hardware technology have coincided with advances in vSphere. Virtual machines can now support up to 1TB RAM and 32 vCPUs, and are capable of running even the largest Exchange mailbox servers. VMware ESXi hosts can now take advantage of up to 160 CPUs and 2TB of physical RAM. Network improvements such as Jumbo Frames and TCP Segment Offload have lowered overall CPU usage. These and other enhancements make vSphere capable of meeting performance requirements for even the most demanding Exchange workloads.

Third-party testing of Exchange Server 2010 in virtual operation has been completed with Microsoft's JetStress and LoadGen tools, the standard tools for Exchange performance analysis. These tests show that performance for a virtualized Exchange server is comparable to a non-virtualized server running on the same hardware. This proved to be true for all Exchange Server 2010 server roles, including the mailbox server. With concerns over relative performance eliminated, many more Exchange administrators are finding the flexibility, enhanced availability, and lower costs associated with virtualization very attractive in supporting an Exchange infrastructure.

3.2 Key Performance Considerations

A variety of factors can affect Exchange Server 2010 performance on vSphere, including processor and memory allocation to the guest virtual machine, storage layout/design, virtual machine placement, and high availability methods, to name a few. The following are some tips for achieving the best possible performance:

- Fully understand your organization's business and technical requirements for implementing Exchange.
- Fully understand the Exchange workload requirements. Current workloads can be measured using the *Microsoft Exchange Server Profile Analyzer* (<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=10559>).
- Size for I/O and not just capacity. Dedicating the appropriate number of spindles (disks) can greatly affect performance of an Exchange virtual machine.
- Use Microsoft sizing and configuration guidelines for the Exchange virtual machines.
- Follow the best practices in Section 2, VMware ESX Host Best Practices for Exchange, to make sure that the ESXi host environment is optimized for enterprise applications such as Exchange.

3.3 Performance Testing

Every Exchange environment is different, with varying business and technical requirements, many server and storage options, and requirements for integrating with third-party software solutions such as anti-virus, anti-spam, PDA messaging, and so on. Due to the many variables, it is strongly recommended that each organization test performance on their particular mix of server, storage, and software to determine the best design for their Exchange environment. In addition, several VMware server and storage partners have performed testing to validate Exchange performance on vSphere. Both of these options are discussed in this section.

3.3.1 Internal Performance Testing

Microsoft provides tools to measure the performance of Microsoft Exchange Server architectures. LoadGen is used to measure performance of both Exchange Server 2007 and Exchange Server 2010. For both versions a storage qualification tool, JetStress, can be used to evaluate the storage configuration.

Because in-guest time is subject to minute fluctuations based on system load, VMware strongly discourages executing any performance analysis tool inside a virtual machine. Accurate measurements are best attained by creating tools that can use the host timer, or by using tests such as LoadGen that contain client/server architectures. For client server tests, the server-under-test may be on a virtual machine while the client is on a native system. This results in accurate measurements at the client.

JetStress, however, does not provide a mechanism for using native timing measurements. The accuracy of JetStress measurements is subject to load on the hypervisor, the processor model, the version of ESXi, and has been known to display a bias. This means that JetStress can report performance metrics that outperform native. Conclusions about the actual performance of this tool in the virtual machine are difficult to find, but it is best to consider JetStress results as broad indicators of performance. Generally speaking, Jetstress results run on newer processors with vSphere 4 or later will show minimal and tolerable error.

3.3.2 Partner Performance Testing

VMware and its OEM partners have been working together for years to characterize Exchange performance. The following table summarizes our performance testing for Exchange 2010 running on vSphere.

Table 2. Performance Testing Summary

Partner	Type	Summary	Resource
VMware	Performance (2010 on vSphere 5)	This paper examines the performance implications of scaling up an Exchange Server 2010 mailbox server virtual machine in vSphere 5 in comparison to vSphere 4.1 and scaling out with multiple Exchange Server 2010 virtual machines .	http://www.vmware.com/files/pdf/exchange-perf-vsphere5.pdf
Dell	Storage Sizing & Performance (2010 on vSphere 4.1)	Provides server and storage sizing best practices for deploying Exchange 2010 on vSphere and Dell PowerEdge blade servers with a Dell Equallogic SAN	http://i.dell.com/sites/content/business/solutions/whitepapers/en/Documents/vmware-vsphere-equallogic.pdf
EMC	Storage Sizing & Performance (2010 on vSphere 4.1)	Presents methodologies and guidelines for designing a scalable Exchange 2010 environment on EMC VNX5700 series unified storage and vSphere.	http://www.emc.com/collateral/hardware/white-papers/h8152-exchange-performance-vnx-wp.pdf
EMC	Storage Sizing & Performance (2010 on vSphere 4.1)	Outlines a solution for Exchange 2010 on vSphere 4.1 powered by EMC VMAX, EMC SRDF and protected with VMware Site Recovery Manager.	http://www.emc.com/collateral/hardware/white-papers/h8151-virtual-exchange-vmax-vmware-wp.pdf
EMC	Storage Sizing & Performance (2010 on vSphere 4.1)	Provides performance results, storage design guidelines and best practices for designing virtualized VMAX storage with Exchange 2010 on vSphere	http://www.emc.com/collateral/solutions/white-papers/h7018-exchange-vmax-vp-vsphere-wp.pdf

3.4 Ongoing Performance Monitoring and Tuning

Traditional Exchange Server performance monitoring leverages the Microsoft Windows performance monitor tool PerfMon to collect statistics. Exchange integrates with PerfMon to provide familiar counters that indicate system performance. But, as with all in-guest measurement tools, time-based performance measurements are subject to error. The degree to which the measurements are inaccurate depends on the total load of the ESXi host. But, generally it is safe to assume the results are no more than 10% in error if CPU utilization stays below 80%.

Exchange administrators should pay close attention to the counters listed in Table 3. Refer to *Performance Monitoring and Analysis* (<http://communities.vmware.com/docs/DOC-3930>) for more information on these counters and their interpretation.

Table 3. Performance Counters of Interest to Exchange Administrators

Subsystem	esxtop Counters	vCenter Counter
CPU	%RDY	Ready (milliseconds in a 20,000 ms window)
	%USED	Usage
Memory	%ACTV	Active
	SWW/s	Swapin Rate
	SWR/s	Swapout Rate
Storage	ACTV	Commands
	DAVG/cmd	Device Latency
	KAVG/cmd	Kernel Latency
Network	MbRX/s	packetsRx
	MbTX/s	packetsTx

This table indicates a few key counters that should be added to the list of inspection points for Exchange administrators. Of the CPU counters, the total used time indicates system load. Ready time indicates overloaded CPU resources. A significant swap rate in the memory counters is a clear indication of a shortage of memory, and high device latencies in the storage section point to an overloaded or misconfigured array. Network traffic is not frequently the cause of most Exchange performance problems except when large amounts of iSCSI storage traffic are using a single network line. Check total throughput on the NICs to see if the network is saturated.

3.5 vSphere Performance Features and Exchange

VMware virtual machines always benefit from newer versions of vmxnet, so we recommend the newest version of vmxnet for Microsoft Exchange virtual machines. The VMware paravirtualized SCSI driver is optimized for high IO environments, but need not be used when a virtual machine generates less than 2000 IOPS. For more information see *VMware KB 1017652* (<http://kb.vmware.com/kb/1017652>).

vSphere 4 also introduced support for second-generation virtualization assist in Intel processors. As with VMware Infrastructure 3, which contained support for similar functionality from AMD, VMware recommends using this additional memory management for most workloads. AMD's memory management, called *Rapid Virtualization Indexing* (RVI), can be leveraged with ESX 3.5 or later. Intel's memory management, called *Extended Page Tables* (EPT), requires ESX 4.0 or later.

4. Exchange 2010 Capacity Planning

4.1 Capacity Planning Process Overview

Sizing of an Exchange 2010 environment is a complex process with many variables including business requirements, anticipated mailbox workloads, and hardware platform, to name a few. The good news is that sizing an Exchange 2010 environment on vSphere is nearly the same as sizing for physical servers.

First, decide whether or not you'll be clustering the mailbox servers.

- If you choose to use standalone mailbox servers protected by HA, use the building block approach defined in Section 4.6.1, The Building Block Approach (Standalone Mailbox Servers).
- If you decide to implement Database Availability Groups (DAGs), use the DAG approach described in Section 4.6.2, The DAG Approach.

Storage sizing and configuration can vary depending on the storage array used and many vendors have unique enhancements to the storage solution that can increase availability, speed recovery, and enhance performance, etc. To optimize performance and take advantage of these features, it is highly recommended that the storage partner be included in the design effort.

There are many facets to an Exchange 2010 deployment besides sizing. Exchange 2010 can be deployed into some very complex, multi-site architectures that should be designed with the assistance of an Exchange expert, whether that person is an internal company resource or a partner with experience deploying both Exchange and vSphere.

When sizing virtual machines, note the following:

- Smaller virtual machines (CPU and RAM) can be moved faster with vMotion than larger virtual machines. For example, a mailbox server virtual machine with 2 vCPU and 9GB RAM can move to another ESXi host with vMotion much quicker than a virtual machine with 4 vCPU and 32GB RAM. Although larger virtual machines can support more users, smaller virtual machines can be more agile in a vSphere environment.
- Size CPU and RAM resources conservatively, and only adjust as required. One of the primary benefits a virtual machine provides is the ability to allocate more CPU or RAM at any time. This means administrators can size their servers based on realistic estimates and only increase resources as required.
- When deciding on the number and size of virtual machine building blocks, consider how this will impact licensing costs for the operating system and applications. Depending on your licensing agreements, more virtual machines may increase licensing costs. Find the right balance between cost and flexibility for your environment.

4.2 Mailbox Server: Planning for CPU

4.2.1 Passive Database Overheads


The Database Availability Group (DAG) feature in Exchange 2010 necessitates a different approach to sizing the Mailbox Server role, forcing the administrator to account for both *active* and *passive* mailboxes. Mailbox Servers that are members of a DAG can host one or more passive databases in addition to any active databases for which they may be responsible.

Each passive database adds an additional *10% to the CPU requirements of the mailbox server hosting the active copy*.

Figure 6 illustrates this principle. We have three Exchange mailbox servers, each with an active database (DB1a denotes database 1 active) and two passive databases from the other two mailbox servers (DB1p denotes database 1 passive). Each passive copy of DB1a requires 10% extra processing on the server hosting DB1a, for a total of 20% extra CPU overhead.

So each mailbox server in this example *requires 20% additional processing power* to account for passive database copies.

Figure 6. Passive Database Overhead



4.2.2 Exchange Server Minimums and Recommended Maximums

It is important to note the new Exchange 2010 minimums and recommended maximums. For example, mailbox servers are now required to have at least two processor cores (or vCPUs). The recommended maximum for a single-role server is 12 cores. vSphere can scale up to 32 vCPUs.

Table 4. Exchange Server Minimums and Recommended Maximums

Exchange 2010 server role	Minimum	Recommended Maximum
Edge Transport	1 x processor core	12 x processor cores
Hub Transport	1 x processor core	12 x processor cores
Client Access	2 x processor core	12 x processor cores
Unified Messaging	2 x processor core	12 x processor cores
Mailbox	2 x processor core	12 x processor cores
Client Access/Hub Transport combo-role (Client Access and Hub Transport roles running on the same physical server)	2 x processor core	12 x processor cores
Multi-role (Client Access, Hub Transport and Mailbox server roles running on the same physical server)	2 x processor cores	24 x processor cores

4.2.3 Megacycles

A *megacycle* is a unit of measurement used to represent processor capacity. To give a rough estimate, a 1GHz processor can produce approximately 1,000 megacycles of CPU throughput. For a larger example, a two-socket, quad-core server (eight cores) with 3.33GHz CPUs can produce approximately 26,400 megacycles.

Each Exchange user placed on the server subtracts from this capacity at varying rates depending on the activity and size of the mailbox. Don't forget that we must take into account CPU requirements for both the active and passive mailboxes that are hosted on the server (see Section 4.1, Capacity Planning Process Overview).

From Microsoft TechNet (<http://technet.microsoft.com/en-us/library/ee712771.aspx>):

Megacycles are estimated based on a measurement of Intel Xeon x5470 3.33GHz processors (2 x 4 core arrangement). A 3.33-GHz processor core = 3,300 megacycles of performance throughput. Other processor configurations can be estimated by comparing this measured platform to server platforms tested by the Standard Performance Evaluation Corporation (SPEC). For details, see the SPEC CPU2006 results at the Standard Performance Evaluation Corporation Web site (<http://www.spec.org/>).

4.2.4 User Profile and Message Activity

Each mailbox (active or passive) has some degree of impact to the overall CPU capacity of the mailbox server, depending on the activity level of the mail user. For example, a mailbox that sends/receives 150 messages a day generates 3 megacycles of CPU activity if the mailbox is in an active database, or .45 megacycles of activity if the mailbox is in a passive database.

Also listed in Table 5 are I/O numbers to aid in storage planning. In the previous example, if we choose to run a standalone mailbox server, we can expect .18 IOPS per mailbox. If we choose to cluster with DAGs that number drops to .15 IOPS, but there is an impact on the passive DAG node.

Table 5. Per Mailbox Database Cache, IOPS, and CPU Estimates Based on User Profile and Message Activity*

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
250	15	0.3	0.25	5	0.75
300	18	0.36	0.3	6	0.9
350	21	0.42	0.35	7	1.05
400	24	0.48	0.4	8	1.2
450	27	0.54	0.45	9	1.35
500	30	0.6	0.5	10	1.5

* <http://technet.microsoft.com/en-us/library/ee712771.aspx>

These metrics, combined with an understanding of the client protocols used to access Exchange resources (Microsoft Outlook, Outlook Anywhere, Outlook Web Access, ActiveSync, and BlackBerry devices, etc.), provide the foundation for planning Exchange CPU, memory, storage, and network requirements. The benefit of profiling Exchange users in this way is that it is platform independent. These same metrics can be used for IBM Lotus Notes, Novell GroupWise, or any other messaging platform to plan resource requirements for an Exchange Server 2010 migration.

If your organization is currently running an older version of Microsoft Exchange, these requirements can be easily determined using the *Microsoft Exchange Server Profile Analyzer* (<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=10559>). This tool can be installed directly on an Exchange server or on an administrator's desktop to collect mailbox usage statistics across the entire Exchange environment, specific Exchange servers, or individual Exchange storage groups and databases as a representative sample of the entire environment.

4.2.5 Designing for Peak Utilization

It is recommended that standalone servers with only the mailbox role be designed to not exceed 70% utilization during peak period. If deploying multiple roles on the server, the mailbox role should be designed not to exceed 35%.

For solutions leveraging mailbox resiliency, it is recommended that the configuration not exceed 80% utilization after a single or double member server failure when the server only has the mailbox role installed. If deploying multiple roles on the server, then the mailbox role should be designed not to exceed 40%.

CPU utilization is determined by taking the CPU Megacycle Requirements and dividing it by the total number of megacycles available on the server (which is based on the CPU and number of cores).

4.3 Mailbox Server: Planning for Memory

4.3.1 Minimum and Recommended Memory

Table 6. Memory Configurations for Exchange 2010 Servers Based on Installed Server Roles*

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 is variable based on the user profile.

* <http://technet.microsoft.com/en-us/library/dd346700.aspx>

4.3.2 Determining Database Cache Size

The first step in planning for Mailbox Server memory is to determine the amount of required database cache by multiplying the mailbox count by the memory requirements based on the user profile. For example, 4,000 users sending/receiving 150 messages per day will require 36GB of database cache. (4000 * 9MB = 36GB).

Table 7. Per Mailbox Database Cache, IOPS, and CPU Estimates Based on User Profile and Message Activity*

Messages sent or received per mailbox per day	Database cache per mailbox in megabytes (MB)
50	3
100	6
150	9
200	12
250	15
300	18
350	21
400	24
450	27
500	30

* <http://technet.microsoft.com/en-us/library/ee712771.aspx>

4.3.3 Determining Total Memory

The next step is to determine the amount of required physical memory by determining which server configuration provides 36GB of database cache. For example, a single role Mailbox server with 48GB of physical RAM will provide 39.2GB of database cache. Therefore, 48GB of physical RAM is the ideal memory configuration based on this mailbox count/user profile.

Table 8. Determining Total Memory

Server Physical Memory (RAM)	Database Cache Size: (Mailbox role only)	Database Cache Size: Multiple-role (for example, Mailbox + Hub Transport)
2GB	512MB	Not supported
4GB	1GB	Not supported
8GB	3.6GB	2GB
16GB	10.4GB	8GB
24GB	17.6GB	14GB
32GB	24.4GB	20GB
48GB	39.2GB	32GB
64GB	53.6GB	44GB
96GB	82.4GB	68GB
128GB	111.2GB	92GB

4.4 Mailbox Server: Planning for Storage

Each Exchange server role has unique storage requirements with respect to throughput (IO) and capacity. Planning storage configurations for the mailbox server role requires knowledge of the existing user profile. Microsoft has defined user profiles by average messages sent and received per day per user. This allows for more accurate planning when migrating from email systems other than Microsoft Exchange. The user profile has a direct impact on overall IO requirements, and knowing these requirements can help you and your storage vendors in designing an optimal storage solution. In addition to the average mail sent and received; mobile devices, archiving solutions, and anti-virus programs should be taken into consideration as contributors to overall IO. The *Microsoft Exchange Server Profile Analyzer* (<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=10559>) can help in gathering much of this data from an existing Exchange environment.

Microsoft has detailed guidelines available on TechNet in the *Mailbox Server Storage Design* (<http://technet.microsoft.com/en-us/library/dd346703.aspx>) section of the Exchange Server 2010 documentation (<http://technet.microsoft.com/en-us/library/dd346703.aspx>). Microsoft also provides the *Exchange 2010 Mailbox Server Role Requirements Calculator* (<http://blogs.technet.com/b/exchange/archive/2009/11/09/3408737.aspx>) to assist in planning the storage design of the Mailbox server role. VMware recommends that you follow Microsoft's best practices along with your storage vendor's best practices to achieve an optimal storage configuration for Exchange Server 2010. For examples on using the calculator, see Section 5, Sizing Examples, and the companion document *Microsoft Exchange 2010 on VMware: Design and Sizing Examples*.

4.5 Hub Transport and Client Access Server Planning

4.5.1 Server Role Ratios by Processor Core

Microsoft provides some general guidelines to plan the number of Hub Transport and Client Access Servers based on processor core ratio to the mailbox server role. These guidelines are meant as a general rule, but your availability requirements, client protocol usage, and unique configuration requirements may indicate more optimal configurations.

Note In Exchange 2010, the CAS server plays a more important role in the architecture. The CAS ratios have changed from Exchange 2007.

The following table summarizes the Microsoft-recommended guidelines.

Table 9. Recommended Server Role Ratios Based on Processor Core*

Server role ratio	Recommended processor core ratio
Mailbox: Hub	7:1 (No antivirus scanning on hub)
	5:1 (With antivirus scanning on hub)
Mailbox: Client Access	4:3
Mailbox: Combined Hub/CAS	1:1

* <http://technet.microsoft.com/en-us/library/dd346701.aspx>

When doing processor core ratios, remember to factor in the *expected peak utilization* of your mailbox servers. For example, Microsoft recommends 70% peak utilization for a standalone mailbox server. If your mailbox server is configured with 6 cores at 70% peak utilization, the mailbox server is actually *using* $6 \times .70$ or 4.2 cores. Use this new number to calculate Hub and CAS core ratios:

- $6 \text{ (total number of mailbox cores)} \times .70 \text{ peak utilization} = 4.2 \text{ (utilized mailbox cores)}$.
- $4.2 \text{ (utilized mailbox cores)} / 5:1 = .84 \text{ Hub Transport cores (with anti-virus scanning)}$.
- $4.2 \text{ (utilized mailbox cores)} / 4:3 = 3.2 \text{ CAS cores}$.

4.5.2 Memory Requirements by Processor Core Count

Microsoft provides guidance for planning the memory requirements for the Hub Transport and Client Access Servers based on the number of processor cores recommended for each server role. The following table summarizes the guidelines that Microsoft recommends.

Table 10. Recommended Server Role Memory Based on Processor Core*

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 is variable based on the user profile.

* <http://technet.microsoft.com/en-us/library/dd346700.aspx>

4.6 Scaling Exchange for the Enterprise

4.6.1 The Building Block Approach (Standalone Mailbox Servers)

The building block approach is a recommended best practice for creating a standalone Exchange Mailbox Servers running on vSphere using pre-sized virtual machine configurations. Exchange servers that have been divided into virtual machine building blocks (as opposed to larger, monolithic Exchange servers) can simplify server sizing during the initial deployment and create a highly scalable solution using virtual machines with predictable performance patterns. Testing by VMware and its partners has focused on four primary sizes for mailbox virtual machine building blocks consisting of 500, 1000, 2000, and 4000 users. These configurations have known performance profiles that can be leveraged for rapid Exchange server sizing as well as easily scaling environments as additional Exchange servers need to be brought online.

The following table presents some pre-sized virtual machine building block examples designed to host mailboxes with an average of 150 messages sent/received per day. The same principles are used for sizing profiles ranging from 50 to 550 messages sent/received per day.

Table 11. Building block CPU and RAM requirements for mailboxes with 150 messages sent/received per day

Building Block	500	1000	2000	4000
Profile	150 sent/received	150 sent/received	150 sent/received	150 sent/received
Megacycle Requirement	1,500	3,000	6,000	12,000
vCPU (based on 3.33GHz processor-based server)	2 (Minimum) (.6 Actual)	2 (Minimum) (1.3 Actual)	4 (2.6 Actual)	6 (5.1 Actual)
Cache Requirement	4.5GB	9GB	18GB	36GB
Total Memory Size	16GB	16GB	24GB	48GB

* Based on information from <http://technet.microsoft.com/en-us/library/ee712771.aspx>.

The sizing process begins with understanding and applying Microsoft guidelines for each server role, as represented by the following high-level processes:

- Design the mailbox server building block.
 - Define *current workloads* using the *Microsoft Exchange Server Profile Analyzer* (<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=10559>).
 - Choose an appropriate (500, 1000, 2000, and 4000 user blocks have been tested and validated, although larger building blocks may be possible).
 - Apply Microsoft guidelines to *determine the CPU requirements*.
 - Apply Microsoft guidelines to *determine the amount of memory required*.
 - Utilize the *Exchange 2010 Mailbox Server Role Requirements Calculator* (<http://blogs.technet.com/b/exchange/archive/2009/11/09/3408737.aspx>) from Microsoft to determine storage requirements.

- Design the peripheral server roles.
 - Determine how many mailbox server building blocks are needed.
 - Calculate the number of mailbox server processor cores.
 - Use Microsoft Guidelines for Server Role Ratios to calculate processor and memory requirements for the Hub Transport roles.
 - Use Microsoft Guidelines for Server Role Ratios to calculate processor and memory requirements for the Client Access Server roles.
- Allocate one or more virtual machines for each server role to satisfy the previously calculated number of processor cores and amount of memory.
- Determine how the virtual machines will be distributed across ESXi hosts.
- Aggregate virtual machine requirements plus some overhead to size each ESXi host. The overhead is important if you want to minimize the performance hit during the loss of one of your ESXi hosts. A typical guideline when choosing the number of required hosts is $n+1$, where n is the number of hosts required to run the workload at peak utilization. $N+1$ allows you to design for the possibility of losing one host from your VMware cluster without taking a huge performance hit during failover.

4.6.2 The DAG Approach

The new Database Availability Group (DAG) feature in Exchange 2010 necessitates a different approach to sizing the Mailbox Server role, forcing the administrator to account for both active and passive mailboxes. Mailbox Servers that are members of a DAG can host one or more passive databases in addition to any active databases for which they may be responsible.

The sizing process begins with understanding and applying Microsoft's guidelines for each server role, as represented by the following high-level processes:

- Design the Mailbox Server DAG nodes.
 - Define current workloads using the *Microsoft Exchange Server Profile Analyzer* (<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=10559>).
 - To greatly simplify the capacity planning process, use the *Exchange 2010 Mailbox Server Role Requirements Calculator* (<http://blogs.technet.com/b/exchange/archive/2009/11/09/3408737.aspx>) to calculate CPU, memory, and storage sizing.
 - Alternatively, if you prefer a manual process:
 - Apply Microsoft guidelines to determine the CPU and memory requirements. Inputs will include, number of mailboxes, mailbox profile, number of servers in the DAG, number of passive database copies, and several other custom parameters.
 - Use the *Exchange 2010 Mailbox Server Role Requirements Calculator* from Microsoft to determine storage requirements.
- Design the peripheral server roles.
 - The *Exchange 2010 Mailbox Server Role Requirements Calculator* will also recommend CPU and memory for the CAS and Hub Transport roles.
 - Alternatively, if you prefer a manual process:
 - Count the number of mailbox server processor cores.
 - Multiply by expected CPU utilization. This should be less than 80% for a clustered mailbox server (for example, $16 \text{ cores} * .80 = 14 \text{ cores}$ (rounded from 12.8)).

- Use the modified number of mailbox cores and Microsoft Guidelines for Server Role Ratios to calculate processor and memory requirements for the Hub Transport roles.
- Use the modified number of mailbox cores and Microsoft Guidelines for Server Role Ratios to calculate processor and memory requirements for the Client Access Server roles.
- Allocate one or more virtual machines for each server role to satisfy the previously calculated number of processor cores and amount of memory.
- Determine how the virtual machines will be distributed across ESXi hosts.
- Aggregate virtual machine requirements plus some overhead to size each ESXi host. The overhead is important if you want to minimize the performance hit during the loss of one of your ESXi hosts. A typical guideline when choosing the number of required hosts is $n+1$, where n is the number of hosts required to run the workload at peak utilization. $N+1$ allows you to design for the possibility of losing one host from your VMware cluster without taking a huge performance hit during failover.

4.7 vSphere Limitations

Although vSphere can support very large virtual machines, there are some limits that you should be aware of when Capacity Planning for Exchange. Below are some of the most common configurations to be aware of. Complete documentation can be found in the *vSphere Configuration Maximums* (<http://www.vmware.com/pdf/vsphere5/r50/vsphere-50-configuration-maximums.pdf>) document on VMware.com.

- vSphere virtual machines are limited to 32 vCPUs and 1TB of RAM. Microsoft recommendations are well below these limitations.
- Each ESXi host can only accommodate up to 256 LUNs. When clustering ESXi hosts this limit effectively extends to the cluster.
- Each virtual disk (VMDK) is limited to 2TB, raw-device mappings and VMFS volumes are limited to 64TB.

In the sizing examples below and in the Design and Sizing Examples document, we've taken the ESXi limitations into account, especially when configuring storage. For example, in the DAG sizing section, we limited database sizes to 1TB to prevent coming too close to our 2TB virtual disk limit.

5. Sizing Examples

Every environment is different and some organizations use email more heavily than others. To accurately determine your mailbox profile requirements, use the *Microsoft Exchange Server Profile Analyzer* (<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=10559>). It is also highly recommended that you work with a partner that is extremely well-versed in Exchange Architecture to design for proper performance in your specific environment.

Note that these examples do not take into account a particular storage solution. Many VMware storage partners have performed extensive testing on building blocks of varying capacities and workloads. See the *Microsoft Exchange 2010 on VMware: Partner Resource Catalog* for storage-specific implementation details.

5.1 Standalone Mailbox Server – 16,000 Users (150 sent/received)

Using the Microsoft sizing guidelines and the building block approach, we apply the formula to size a 16,000-user environment with mailbox profiles measured at 150 messages sent/received per user per day. The mailboxes are distributed across four mailbox server building blocks of 4,000 users each. The following calculations are meant to serve as an example of the sizing process—every organization's implementation details will vary.

In our example, we use the following mailbox profile:

- 150 messages sent/received per day.
- Average message size of 75 KB.
- 2048MB mailbox quota.

5.1.1 Calculate Mailbox Server CPU Requirements (4,000 Users/VM)

Microsoft guidelines recommend using megacycles to determine the amount of CPU required by the mailbox server. Each user consumes resources on the processor at varying rates depending on message activity. Because we are not implementing Database Availability Groups in this example, we can use the guidance for standalone mailbox servers from Table 12.

Table 12. Mailbox Server CPU Recommendations*

Messages sent or received per mailbox per day	Megacycles for stand-alone mailbox
50	1
100	2
150	3
200	4
250	5
300	6
350	7
400	8
450	9

* <http://technet.microsoft.com/en-us/library/ee712771.aspx>

The first step in sizing our 4,000-user building block is to apply these Microsoft guidelines to determine the CPU requirements.

Example:

Your organization plans to support 16,000 users with 4,000 mailboxes per Exchange Mailbox Server and a 2GB mailbox quota. You have used the Exchange Server Profile Analyzer to determine that your users average 150 messages sent and received per day, with an average message size of 75 KB. The standard hardware build includes 16 core servers (4x4) with each CPU at 3.33GHz.

- 16,000 mailboxes/4,000 mailboxes per server = *4 mailbox servers*.
- 4,000 mailboxes per server * 3 megacycles per active mailbox = *12,000 megacycles required*.
- Each 3.33GHz processor core can produce *3,330 megacycles*.
- Microsoft recommends *70% peak utilization* of a standalone mailbox server.
- 12,000 megacycles required/.70 = *17,143 adjusted megacycles required*.
- 17,143 megacycles/3,330 = *6 processor cores per server (rounded up from 5.15 actual)*.
- 4 mailbox servers x 6 processor cores per server = *24 processor cores*.
- Our mailbox servers will utilize 60% of their processing capability at peak utilization. Adjusting the number of virtual processors to 5 would help increase processor utilization.

5.1.2 Calculate Mailbox Server Memory Requirements (4,000 Users/VM)

Use the first table to determine the required amount of database cache and the second table to pick the total server memory that will satisfy the database cache requirement.

Table 13. Mailbox Server Database Cache Requirements*

Messages sent or received per mailbox per day	Database cache per mailbox in megabytes (MB)
50	3
100	6
150	9
200	12
250	15
300	18
350	21
400	24
450	27
500	30

* <http://technet.microsoft.com/en-us/library/ee712771.aspx>

Table 14. Default Mailbox Database Cache Sizes*

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

128GB	111.2GB	92GB
-------	---------	------

* <http://technet.microsoft.com/en-us/library/ee832793.aspx>

Example:

Given our previous example of 16,000 mailboxes with 4,000 mailboxes per server with each mailbox sending/receiving 150 messages per day:

- 9MB x 4,000 mailboxes = 36GB required database cache.
- According to Table 14 we need 48GB of total memory in the mailbox server to provide 39.2GB of database cache. This will satisfy the 36GB requirement.

5.1.3 Calculate Mailbox Server Storage Requirements

Table 15. Environment Configuration

Exchange Environment Configuration	Value
Global Catalog Server Architecture	64-bit
Server Multi-Role Configuration (MBX+CAS+HT)	No
High Availability Deployment	No
Site Resiliency Deployment	No
Number of Mailbox Servers	4
Global Catalog Server Architecture	64-bit
Server Multi-Role Configuration (MBX+CAS+HT)	No

Table 16. Data Configuration

Exchange Data Configuration	Value
Data Overhead Factor	20%
Mailbox Moves/Week Percentage	1%
Dedicated Maintenance/Restore LUN?	Yes
LUN Free Space Percentage	20%

Table 17. Mailbox Configuration

Tier-1 User Mailbox Configuration	Value
Total Number of Tier-1 User Mailboxes	16000
Projected Mailbox Number Growth Percentage	0%
Total Send/Receive Capability/Mailbox/Day	150 messages
Average Message Size (KB)	75
Mailbox Size Limit (MB)	2048
Personal Archive Mailbox Size Limit (MB)	0
Deleted Item Retention Window (Days)	14
Single Item Recovery	Enabled
Calendar Version Storage	Enabled
IOPS Multiplication Factor	0.00
Desktop Search Engines Enabled (for Online Mode Clients)	No
Predict IOPS Value?	Yes

Table 18. Backup Configuration

Backup Configuration	Value
Backup Methodology	Software VSS Backup/Restore
Backup Frequency	Weekly Full/Daily Incremental
Database and Log Isolation Configured	No
Backup/Truncation Failure Tolerance	3
Network Failure Tolerance (Days)	0

Table 19. Storage Configuration

Disk Configuration	Disk Capacity	Disk Type
Database	300GB	10K RPM FC/SCSI/SAS 3.5"
Log	300GB	10K RPM FC/SCSI/SAS 3.5"
Restore LUN	300GB	10K RPM FC/SCSI/SAS 3.5"

Table 20. Processor Configuration

Server Configuration	Processor Cores/Server	Megacycles/Core
Mailbox Servers	6	3330

5.1.3.1. Mailbox Server Role Storage Requirements Calculator v6.3 Results

Using the Exchange Server 2010 Mailbox Server Role Storage Requirements Calculator version 6.3 and the above variables, our mailbox server storage configuration is summarized as shown in the following table.

Table 21. Mailbox Server Storage Results

Output	Value
Database/log configuration (per server)	58 databases per server 190GB database size + overhead 9GB log size + overhead
Database LUN design (per server)	19 LUNs recommended (9 DB, 9 Log, 1 Restore) 7 databases per LUN DB1-DB7: 1833GB DB LUN/80GB Log LUN DB8-DB14: 1833GB DB LUN/80GB Log LUN DB15-DB21: 1833GB DB LUN/80GB Log LUN DB22-DB28: 1833GB DB LUN/80GB Log LUN DB29-DB35: 1833GB DB LUN/80GB Log LUN DB36-DB42: 1833GB DB LUN/80GB Log LUN DB43-DB49: 1833GB DB LUN/80GB Log LUN DB50-DB56: 1833GB DB LUN/80GB Log LUN DB57-DB58: 524GB DB LUN/23GB Log LUN Restore LUN: 1747GB

RAID configuration (per server)	Databases Disks per Server (RAID 1/0) -110 x 300GB/10K RPM FC/SCSI/SAS 3.5"
	Log Disks per Server (RAID 1/0) -6 x 300GB/10K RPM FC/SCSI/SAS 3.5"
	Restore LUN Disks per Server (RAID 5) -12 x 300GB/10K RPM FC/SCSI/SAS 3.5"

Example:

With a standalone mailbox server, Microsoft recommends placing logs on separate LUNs than their corresponding database files. Given our previous example of 16,000 mailboxes with 4,000 mailboxes per server with each mailbox sending/receiving 150 messages per day:

- 58 databases per server.
- 19 LUNs recommended per server (9 DB, 9 Log, 1 Restore).
- 7 databases per LUN.
- Physical Disk Requirements (per Server):
 - Databases = 110 x 300GB/10K RPM FC/SCSI/SAS 3.5".
 - Log = 6 x 300GB/10K RPM FC/SCSI/SAS 3.5".
 - Restore LUN = 12 x 300GB/10K RPM FC/SCSI/SAS 3.5".

5.1.4 Mailbox Server Summary (4,000-User/150 sent/received)

The following table summarizes the resource requirements for our 4,000-user building block.

Table 22. 4,000-User Building Block Requirements (150 sent/received)

Exchange Role	Physical Resources (per server)
Mailbox server (4 nodes)	CPU: 6 cores Memory: 48GB OS and Application File Storage: 64GB (OS & application files) Database Storage 110 x 300GB/10K RPM FC/SCSI/SAS 3.5" Log Storage 6 x 300GB/10K RPM FC/SCSI/SAS 3.5" Restore LUN Storage 12 x 300GB/10K RPM FC/SCSI/SAS 3.5" Network: 1Gbps

5.1.5 Guest Virtual Machine Configuration

In Table 23 the resource requirements for Exchange are translated into the virtual machine configuration for our 4,000-user mailbox server building block.

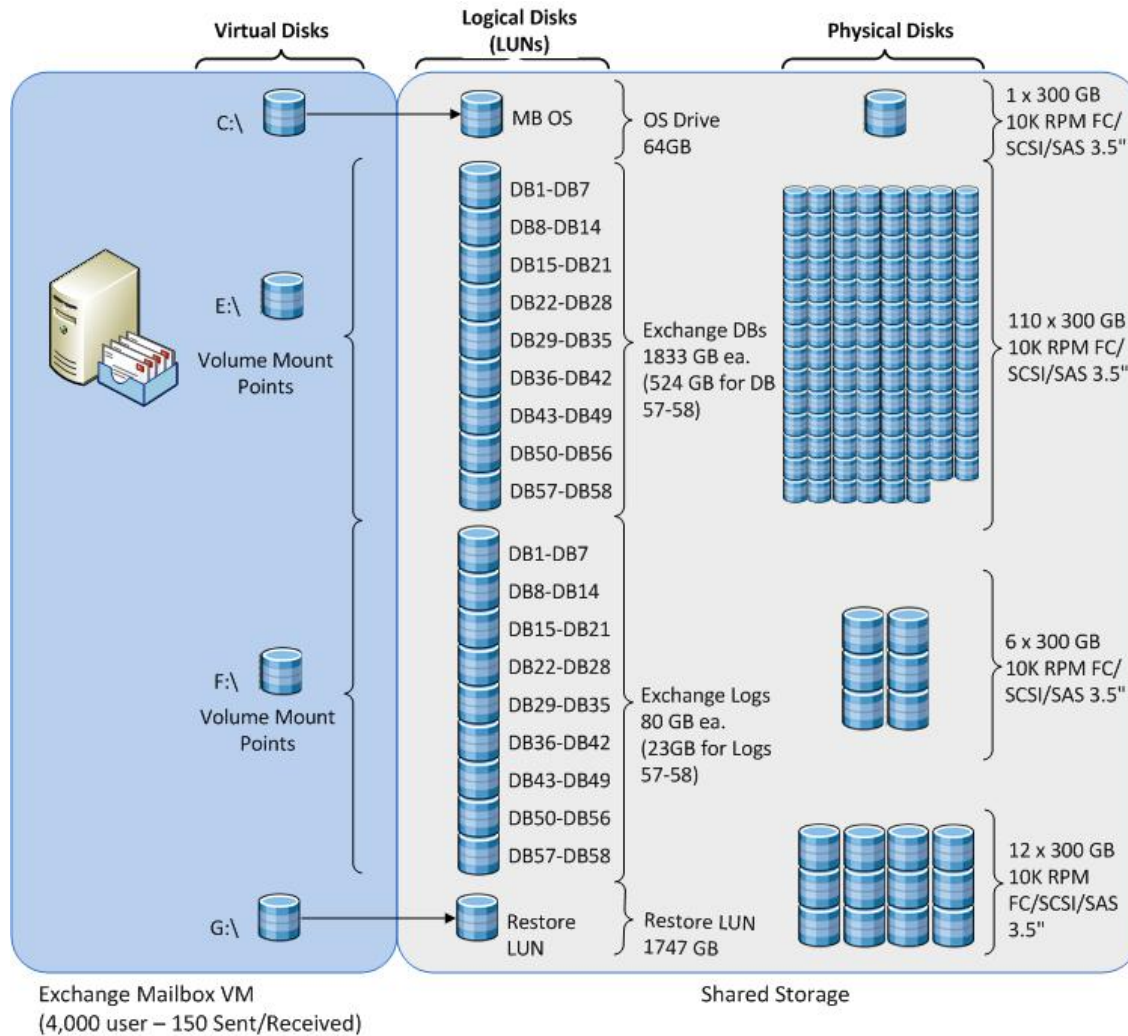
Table 23. Exchange Virtual Machine Configuration

Exchange Role	Virtual Hardware (per VM)
Mailbox Server	<p>CPU: 6 vCPU</p> <p>Memory: 48GB</p> <p>Storage: SCSI Controller 0</p> <p>HDD 1: 64GB (OS & application files)</p> <p>Storage: SCSI Controller 1</p> <p>HDD 2: 1833GB (DB1-DB7 databases)</p> <p>HDD 3: 1833GB (DB8-DB14 databases)</p> <p>HDD 4: 1833GB (DB15-DB21 databases)</p> <p>HDD 5: 1833GB (DB22-DB28 databases)</p> <p>HDD 6: 1833GB (DB29-DB35 databases)</p> <p>HDD 7: 1833GB (DB36-DB42 databases)</p> <p>HDD 8: 1833GB (DB43-DB49 databases)</p> <p>HDD 9: 1833GB (DB50-DB56 databases)</p> <p>HDD 10: 524GB (DB57-DB58 databases)</p> <p>Storage: SCSI Controller 2</p> <p>HDD 11: 80GB (DB1-DB7 logs)</p> <p>HDD 12: 80GB (DB8-DB14 logs)</p> <p>HDD 13: 80GB (DB15-DB21 logs)</p> <p>HDD 14: 80GB (DB22-DB28 logs)</p> <p>HDD 15: 80GB (DB29-DB35 logs)</p> <p>HDD 16: 80GB (DB36-DB42 logs)</p> <p>HDD 17: 80GB (DB43-DB49 logs)</p> <p>HDD 18: 80GB (DB50-DB56 logs)</p> <p>HDD 19: 80GB (DB57-DB58 logs)</p> <p>Storage: SCSI Controller 3</p> <p>HDD 20: 1747GB (Restore LUN)</p> <p>Network: NIC 1</p>

5.1.6 Guest Virtual Machine Storage Interaction

Figure 7 shows how the building block virtual machine interacts with the shared storage.

Figure 7. Building Block Virtual Machine Interaction with Shared Storage



5.1.7 Hub Transport and Client Access Server Role Calculations

Table 24. Recommended Server Role Ratios Based on Processor Core*

Server role ratio	Recommended processor core ratio
Mailbox: Hub	7:1 (No antivirus scanning on hub) 5:1 (With antivirus scanning on hub)
Mailbox: Client Access	4:3
Mailbox: Combined Hub/CAS	1:1

* <http://technet.microsoft.com/en-us/library/dd346701.aspx>

Table 25. Recommended Server Role Ratios Based on Processor Core*

Exchange 2010 server role	Minimum Supported	Recommended
Hub Transport	4GB	1GB per core (4GB minimum)
Client Access	4GB	2GB per core (8GB minimum)
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)

* <http://technet.microsoft.com/en-us/library/dd346700.aspx>

Example:

Given our previous example of 16,000 mailboxes with 4,000 average profile mailboxes per server:

Hub Transport Calculations

- 24 mailbox processor cores * .60 peak utilization = *14.4 mailbox processor cores*.
- 14.4 mailbox processor cores/1:5 hub core ratio with AV = *4 hub processor cores* (rounded for even distribution from 2.88).

Based on your business requirements, you may decide to deploy *2 Hub Transport virtual machines at 2 vCPU per virtual machine*.

To calculate memory for each VM, multiply the number of vCPU * the recommended memory per core from the above table: 2 processor cores x 1GB per core = 2GB RAM; however, we must allocate 4GB of RAM based on the minimum supported configuration.

Client Access Server Calculations

- 24 mailbox processor cores * .60 peak utilization = *14.4 utilized processor cores*.
- 14.4 mailbox processor cores/4:3 Client Access core ratio = *12 client access processor cores* (rounded for even distribution from 10.8).

Based on your business requirements, you may decide to deploy *3 Client Access virtual machines at 4 vCPU per virtual machine*.

To calculate memory for each virtual machine, multiply the number of vCPU * the recommended memory per core from the above table: 4 processor cores x 2GB per core = *8GB RAM*.

5.1.8 Resource Requirements by Server Role

With any application, first plan its resource requirements, then determine the underlying infrastructure requirements to meet the application's needs.

Table 26. Example Exchange Server Role Resource Requirements

Exchange Role	Physical Resources (per server)
Mailbox Server (4 servers)	CPU: 6 cores (60% max utilization) Memory: 48GB OS and Application File Storage: 64GB (OS and application files) DB Storage: 110 x 300GB 10K RPM FC/SCSI/SAS 3.5" (RAID 1/0) Log Storage: 6 x 300GB 10K RPM FC/SCSI/SAS 3.5" (RAID 1/0) Restore LUN: 12 x 300GB 10K RPM FC/SCSI/SAS 3.5" (RAID 5) Network: 1Gbps
Client Access Server (3 servers)	CPU: 4 cores Memory: 8GB Storage: 24GB (OS and application files) Network: 1Gbps
Hub Transport Server (2 servers)	CPU: 2 cores Memory: 4GB Storage: 20GB (OS, application, and log files) 32GB (DB, protocol/tracking logs, and temp files) Network: 1Gbps

5.1.9 Guest Virtual Machine Configuration

The next step is to plan the individual virtual machine configuration based on the required resources for each Exchange server role.

Table 27. Example Exchange Virtual Machine Configuration

Exchange Role	Virtual Hardware (per VM)
Mailbox Server (4 VMs)	CPU: 6 vCPU Memory: 48GB Storage: SCSI Controller 0 HDD 1: 64GB (OS & application files) Storage: SCSI Controller 1 HDD 2: 1833GB (DB1-DB7 databases) HDD 3: 1833GB (DB8-DB14 databases) HDD 4: 1833GB (DB15-DB21 databases) HDD 5: 1833GB (DB22-DB28 databases) HDD 6: 1833GB (DB29-DB35 databases) HDD 7: 1833GB (DB36-DB42 databases) HDD 8: 1833GB (DB43-DB49 databases) HDD 9: 1833GB (DB50-DB56 databases) HDD 10: 524GB (DB57-DB58 databases) Storage: SCSI Controller 2 HDD 11: 80GB (DB1-DB7 logs) HDD 12: 80GB (DB8-DB14 logs) HDD 13: 80GB (DB15-DB21 logs) HDD 14: 80GB (DB22-DB28 logs) HDD 15: 80GB (DB29-DB35 logs) HDD 16: 80GB (DB36-DB42 logs) HDD 17: 80GB (DB43-DB49 logs) HDD 18: 80GB (DB50-DB56 logs) HDD 19: 80GB (DB57-DB58 logs) Storage: SCSI Controller 3 HDD 20: 1747GB (Restore LUN) Network: NIC 1
Client Access Server (3 VMs)	CPU: 4 vCPU Memory: 8GB Storage: SCSI Controller 0 HDD 1: 24GB (OS and application files) Network: NIC 1

Hub Transport Server (2 VMs)	CPU: 2 vCPU Memory: 4GB Storage: SCSI Controller 0 HDD 1: 20GB (OS and application files) Storage: SCSI Controller 0 HDD 2: 32GB (DB, protocol/tracking logs, temp files) Network: NIC 1
------------------------------	---

5.1.10 Planning the Host Hardware Configuration

5.1.10.1. Virtual Machine Distribution

In this example, we have chosen to use four ESXi hosts connected to shared storage to use advanced VMware features such as HA and DRS.

To build infrastructure availability into the architecture, we distribute the nine total virtual machines across four physical ESXi hosts. Initial placement of virtual machines is relatively unimportant, especially if you're utilizing DRS. Given that we're not running a DAG, HA and DRS can be used and are fully supported for the mailbox servers.

Table 28. Example Exchange Virtual Machine Distribution

ESXi Host	Virtual Machines
ESXi Host 1	Exchange Mailbox VM 1 (6 vCPU/48GB RAM) Exchange Client Access VM 1 (4 vCPU/8GB RAM) Exchange Hub Transport VM 1 (2 vCPU/4GB RAM)
ESXi Host 2	Exchange Mailbox VM 2 (6 vCPU/48GB RAM) Exchange Client Access VM 2 (4 vCPU/8GB RAM) Exchange Hub Transport VM 2 (2 vCPU/4GB RAM)
ESXi Host 3	Exchange Mailbox VM 3 (6 vCPU/48GB RAM) Exchange Client Access VM 3 (4 vCPU/8GB RAM)
ESXi Host 4	Exchange Mailbox VM 4 (6 vCPU/48GB RAM)

5.1.10.2. ESXi Host Specifications

Each ESXi host should provide enough physical hardware resources to accommodate the planned workload and provide some headroom in the event of a VMware HA failover or planned vMotion migration of live virtual machines for host hardware maintenance. The following table summarizes the ESXi host hardware configuration based on our example architecture.

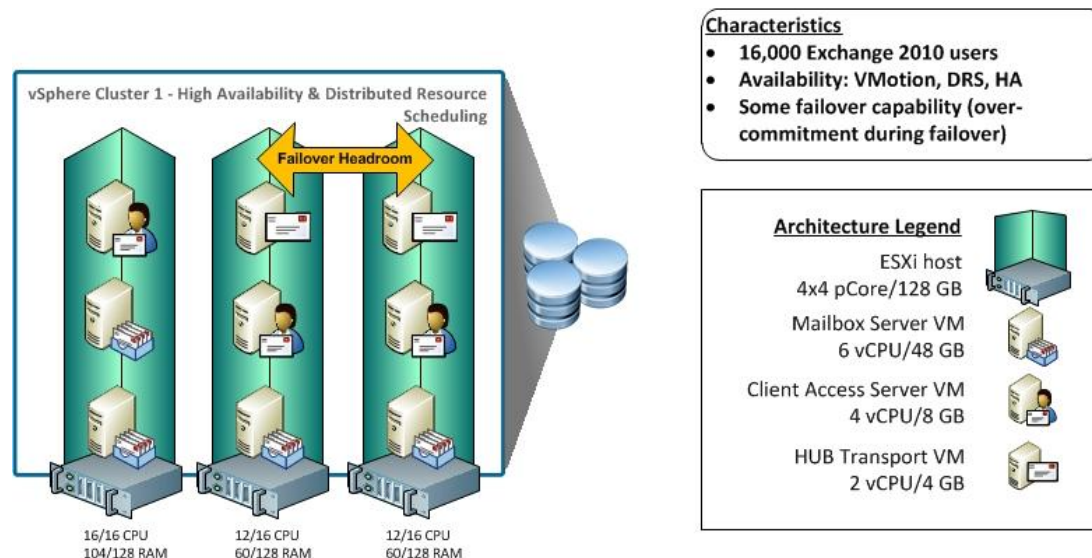
Table 29. Example ESXi Host Hardware Configuration Table

ESXi Host	Virtual Machines
ALL ESXi hosts	16 cores (4x4)
	128GB RAM (extra 14GB above requirements for use in failover)
	2 Fibre Channel HBAs
	4 Gigabit network adapters

5.1.10.3. Initial VM Placement

Although the workloads migrate automatically with DRS (including the mailbox servers, because we're not clustering), Figure 8 is a useful planning tool for initial placement of virtual machines and calculating host failover capacity. At initial placement, ESXi hosts 2 and 3 have the most failover headroom.

Figure 8. Initial Virtual Machine Placement



5.2 Clustered Mailbox Server – 16,000 Users (150 sent/received)

The following example demonstrates the mailbox configuration needed to support 16,000 active users protected by DAG clustering. We use the mailbox calculations to estimate the amount of processing and memory needed for the CAS and Hub Transport Roles. We then translate the estimated resources into virtual machine and host configurations.

In our example, we use the following mailbox profile:

- 150 messages sent/received per day.
- Average message size of 75 KB.
- 2048MB mailbox quota.

In addition, because we are clustering the mailbox role using Database Availability Groups, we need to change our approach to sizing processor, memory, and storage. More resources are needed to support passive database copies).

5.2.1 Calculate Mailbox Server CPU Requirements (4,000 active users)

Microsoft guidelines recommend using megacycles to determine the amount of CPU required by the mailbox server (see Section 4.2.3, Megacycles, for details). Each user will consume resources on the processor at varying rates depending on message activity. Because we are implementing Database Availability Groups in this example, we should use the guidance for both active and passive mailboxes from the following table.

Table 30. Mailbox Server CPU Recommendations*

Messages sent or received per mailbox per day	Megacycles for active mailbox or stand-alone mailbox	Megacycles for passive mailbox
50	1	0.15
100	2	0.3
150	3	0.45
200	4	0.6
250	5	0.75
300	6	0.9
350	7	1.05
400	8	1.2
450	9	1.35
500	10	1.5

* <http://technet.microsoft.com/en-us/library/ee712771.aspx>

The first step in sizing our mailbox server is to apply these Microsoft guidelines to determine the CPU requirements. Without a deep understanding of how DAGs work, we recommend using the Exchange 2010 Mailbox Server Role Requirements Calculator

(<http://blogs.technet.com/b/exchange/archive/2009/11/09/3408737.aspx>) to calculate processor, memory, and storage. In fact, our example here is based on output from the calculator which has been picked apart to demonstrate the basic principles.

Example:

Your organization plans to support 16,000 users with a 2GB mailbox quota. Mailbox servers will be configured in a 4-node DAG with two copies of the data (including the active copy). You have used the Exchange Server Profile Analyzer to determine that your users average 150 messages sent and received per day, with an average message size of 75 KB. The standard hardware build includes 16 core servers (4x4) with each CPU at 3.33GHz.

For this example, we limit each mailbox database to 1TB to keep the number of required LUNs to a minimum. vSphere is limited to 256 LUNs and 2TB per virtual disk. Setting the maximum database size to 1TB also provides enough room to grow a bit. If we want to take advantage of larger database sizes we could deploy raw-device mappings for the data drives. This would allow up to 64TB volumes. According to the storage calculator (see below), setting this 1TB limit results in about 333 mailboxes per database.

Because we are doing a 4-node DAG, during failover of a single node, each remaining mailbox server potentially needs to run four additional databases (approx. 333 mailboxes each). As you may potentially run 5,333 active mailboxes during a failover (4,000 active/1,333 passive), we need to size each server based on all 5,333 mailboxes being active.

$$5,333 \text{ potentially active mailboxes per server} \times 3 \text{ megacycles} = 16,000 \text{ megacycles required}$$

Next, we need to account for passive mailbox overhead. Allocate 10% extra CPU for each passive copy. We can multiply by 1.1 to get our number:

$$16,000 \text{ megacycles} \times 1.1 \text{ (10\% overhead)} = 17,600 \text{ megacycles}$$

Not all passive mailboxes on the server have the potential to be active. In this configuration, each mailbox server is running 12 active and 12 passive databases. If a single node fails, then four of those databases would become active, which we've already accounted for. Now we need to calculate the impact of the remaining passive databases:

- $2,666 \text{ passive mailboxes per server} \times .45 \text{ megacycles} = 1,200 \text{ megacycles required.}$
- $17,600 \text{ active} + 1,200 \text{ passive} = 18,800 \text{ total megacycles required per server.}$

Microsoft recommends 80% peak utilization of a DAG node mailbox server.

- $18,800 \text{ megacycles required} / .80 = 23,500 \text{ adjusted megacycles required.}$
- $23,500 \text{ megacycles} / 3,330 = 8 \text{ processor cores per server (rounded up from 7.05 actual).}$

With eight processor cores per mailbox server you can expect a peak utilization of 71% during failover. Normal operating utilization should be considerably less.

5.2.2 Calculate Mailbox Server Memory Requirements (4,000 Active Users)

Use the first table to determine the required amount of database cache and the second table to pick the total server memory that will satisfy the database cache requirement. You can modify the default database cache size values by making changes to the `msExchESEParamCacheSizeMax` and `msExchESEParamCacheSizeMin` attributes in Active Directory. For additional details, see [How to modify the Store Database maximum cache size in Exchange 2000 Server](http://support.microsoft.com/?kbid=266768) (<http://support.microsoft.com/?kbid=266768>). Use 32KB pages for the cache sizing calculations.

Table 31. Mailbox Server Database Cache Requirements*

Messages sent or received per mailbox per day	Database cache per mailbox in megabytes (MB)
50	3
100	6
150	9
200	12
250	15
300	18
350	21
400	24
450	27
500	30

* <http://technet.microsoft.com/en-us/library/ee832793.aspx>

Table 32. Default mailbox database cache sizes*

Server physical memory (RAM)	Database cache size: (Mailbox role only)	Database cache size: Multiple-role (for example, Mailbox + Hub Transport)
2GB	512MB	Not supported
4GB	1GB	Not supported
8GB	3.6GB	2GB
16GB	10.4GB	8GB
24GB	17.6GB	14GB
32GB	24.4GB	20GB
48GB	39.2GB	32GB
64GB	53.6GB	44GB
96GB	82.4GB	68GB
128GB	111.2GB	92GB

* <http://technet.microsoft.com/en-us/library/ee832793.aspx>

Example:

Given our previous example of 16,000 mailboxes with 4,000 active and 1,333 passive (but potentially active) per server with each mailbox sending/receiving 150 messages per day:

- 9MB x 5,333 mailboxes = 48GB required database cache.
- According to the default mailbox database cache sizes table, 64GB of total memory is needed in the mailbox server to provide 53.6GB of database cache. This will satisfy the 48GB requirement.

5.2.3 Calculate Mailbox Server Storage Requirements

5.2.3.1. Mailbox Server Role Storage Requirements Calculator v6.3 Inputs

Table 33. Environment Configuration

Exchange Environment Configuration	Value
Global Catalog Server Architecture	64-bit
Server Multi-Role Configuration (MBX+CAS+HT)	No
High Availability Deployment	Yes
Site Resiliency Deployment	No
Number of Mailbox Servers	4
Global Catalog Server Architecture	64-bit
Server Multi-Role Configuration (MBX+CAS+HT)	No

Table 34. Data Configuration

Exchange Data Configuration	Value
Data Overhead Factor	20%
Mailbox Moves/Week Percentage	1%
Dedicated Maintenance/Restore LUN?	Yes
LUN Free Space Percentage	20%
Log Shipping Network Compression	Enabled
Log Shipping Compression Percentage	30%

Table 35. Database Copy Configuration

Mailbox Database Copy Configuration	Value
Total Number of HA Database Copy Instances (Includes Active Copy) within DAG	2
Total Number of Lagged Database Copy Instances within DAG	0
Number of HA Database Copy Instances Deployed in Secondary Data Center	0
Number of Lagged Database Copy Instances in Secondary Data Center	0

Table 36. Database Configuration

Database Configuration	Value
Maximum Database Size Configuration	Custom
Maximum Database Size (GB)	1,000
Automatically Calculate Number of Unique Databases/DAG	Yes
Custom Number of Databases/DAG	0

Table 37. Mailbox Configuration

Tier-1 User Mailbox Configuration	Value
Total Number of Tier-1 User Mailboxes	16000
Projected Mailbox Number Growth Percentage	0%
Total Send/Receive Capability/Mailbox/Day	150 messages
Average Message Size (KB)	75
Mailbox Size Limit (MB)	2048
Personal Archive Mailbox Size Limit (MB)	0
Deleted Item Retention Window (Days)	14
Single Item Recovery	Enabled
Calendar Version Storage	Enabled
IOPS Multiplication Factor	0.00

Desktop Search Engines Enabled (for Online Mode Clients)	No
Predict IOPS Value?	Yes

Table 38. Backup Configuration

Backup Configuration	Value
Backup Methodology	Exchange Native Data Protection
Database and Log Isolation Configured	No
Backup/Truncation Failure Tolerance	3
Network Failure Tolerance (Days)	0

Table 39. Storage Configuration

Disk Configuration	Disk Capacity	Disk Type
Database	300GB	10K RPM FC/SCSI/SAS 3.5"
Log	300GB	10K RPM FC/SCSI/SAS 3.5"
Restore LUN	300GB	10K RPM FC/SCSI/SAS 3.5"

Table 40. Processor Configuration

Server Configuration	Processor Cores/Server	Megacycles/ Core
Mailbox Servers	8	3330

5.2.3.2. Mailbox Server Role Storage Requirements Calculator v6.3 Results

Using the *Exchange Server 2010 Mailbox Server Role Storage Requirements Calculator* version 6.3 and the above variables, our Mailbox server storage configuration is summarized in Table 41.

Table 41. Mailbox Server Storage Results

Output	Value
Database/log configuration (per server)	24 databases per server 921GB database size + overhead 44GB log size + overhead
Database LUN design (per server)	25 LUNs recommended (logs will share DB LUNs) 1 database per LUN DB1: 1321GB DB/Log LUN DB2: 1321GB DB/Log LUN DB3: 1321GB DB/Log LUN DB4: 1321GB DB/Log LUN DB5: 1321GB DB/Log LUN DB6: 1321GB DB/Log LUN DB7: 1321GB DB/Log LUN DB8: 1321GB DB/Log LUN DB9: 1321GB DB/Log LUN DB10: 1321GB DB/Log LUN DB11: 1321GB DB/Log LUN DB12: 1321GB DB/Log LUN DB13: 1321GB DB/Log LUN DB14: 1321GB DB/Log LUN DB15: 1321GB DB/Log LUN DB16: 1321GB DB/Log LUN DB17: 1321GB DB/Log LUN DB18: 1321GB DB/Log LUN DB19: 1321GB DB/Log LUN DB20: 1321GB DB/Log LUN DB21: 1321GB DB/Log LUN DB22: 1321GB DB/Log LUN DB23: 1321GB DB/Log LUN DB24: 1321GB DB/Log LUN Restore LUN: 1206GB
RAID configuration (per server)	Database/Log Disks per Server (RAID 5) -138 x 300GB/10K RPM FC/SCSI/SAS 3.5" Restore LUN Disks per Server (RAID 5) -9 x 300GB/10K RPM FC/SCSI/SAS 3.5"

Example:

With a standalone mailbox server, Microsoft recommends placing logs on separate LUNs than their corresponding database files. Because we are replicating the data throughout the DAG, we can place the logs *on the same LUN as the database*, reducing the number of LUNs required.

Given our previous example of 16,000 mailboxes with 4,000 mailboxes per server with each mailbox sending/receiving 150 messages per day:

- 24 databases per server.
- 24 LUNs recommended per server (logs will share DB LUNs).
- 1 database per LUN.
- Physical Disk Requirements (per server):
 - Databases and Logs = 138 x 300GB/10K RPM FC/SCSI/SAS 3.5".
 - Restore LUN = 9 x 300GB/10K RPM FC/SCSI/SAS 3.5".

Note Dedicate a Restore LUN to avoid the storage calculator doubling database LUN sizes.

5.2.4 Mailbox Server Summary (16,000 Active Users/150 sent/received)

The following table summarizes the resource requirements for our DAG nodes.

Table 42. 16,000 Active Mailboxes (150 sent/received) DAG Node Requirements

Exchange Role	Physical Resources (per server)
Mailbox Server (4 nodes)	<p>CPU: 8 cores (82% max utilization)</p> <p>Memory: 64GB</p> <p>OS and Application File Storage: 80GB (OS & Application files)</p> <p>Database and Log Storage 138 x 300GB 10K RPM FC/SCSI/SAS 3.5"</p> <p>Restore LUN Storage 9 x 300GB/10K RPM FC/SCSI/SAS 3.5"</p> <p>Network: 1Gbps</p>

5.2.5 Guest Virtual Machine Configuration

The resource requirements from Table 42 are translated below into virtual machine resources.

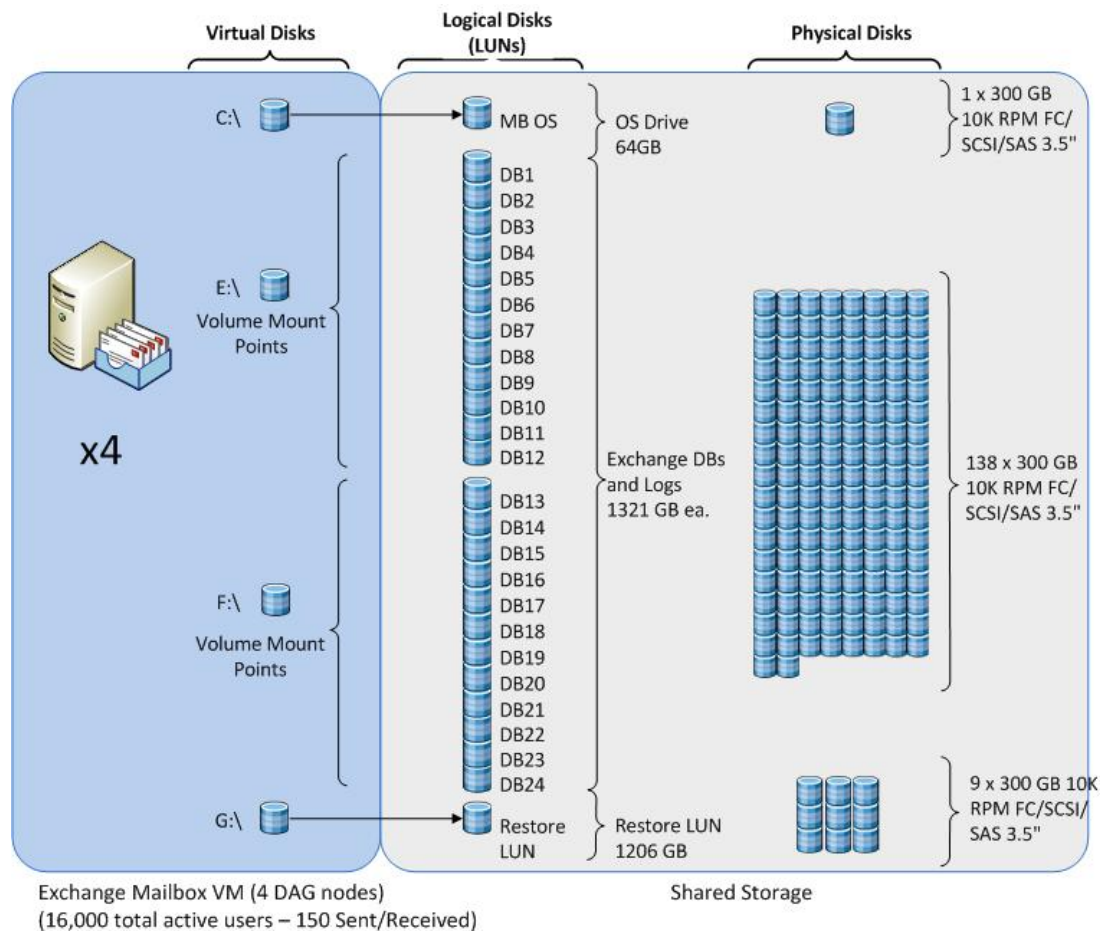
Table 43. Exchange Virtual Machine Configuration

Exchange Role	Virtual Hardware (per VM)
Mailbox Server (4 servers)	CPU: 8 vCPU Memory: 64GB Storage: SCSI Controller 0 HDD 1: 80GB (OS & application files) Storage: SCSI Controller 1 HDD 2: 1321GB (DB1) HDD 3: 1321GB (DB2) HDD 4: 1321GB (DB3) HDD 5: 1321GB (DB4) HDD 6: 1321GB (DB5) HDD 7: 1321GB (DB6) HDD 8: 1321GB (DB7) HDD 9: 1321GB (DB8) HDD 10: 1321GB (DB9) HDD 11: 1321GB (DB10) HDD 12: 1321GB (DB11) HDD 13: 1321GB (DB12) Storage: SCSI Controller 2 HDD 14: 1321GB (DB13) HDD 15: 1321GB (DB14) HDD 16: 1321GB (DB15) HDD 17: 1321GB (DB16) HDD 18: 1321GB (DB17) HDD 19: 1321GB (DB18) HDD 20: 1321GB (DB19) HDD 21: 1321GB (DB20) HDD 22: 1321GB (DB21) HDD 23: 1321GB (DB22) HDD 24: 1321GB (DB23) HDD 25: 1321GB (DB24) Storage: SCSI Controller 3 HDD 26: 1206GB (Restore LUN) Network: NIC 1

5.2.6 Guest Virtual Machine Storage Interaction

Figure 9 shows how each mailbox server virtual machine interacts with the shared storage.

Figure 9. Mailbox Server Virtual Machine Interaction with Shared Storage



5.2.7 Hub Transport and Client Access Server Role Calculations

Table 44. Recommended Server Role Ratios Based on Processor Core*

Server role ratio	Recommended processor core ratio
Mailbox: Hub	7:1 (No antivirus scanning on hub)
	5:1 (With antivirus scanning on hub)
Mailbox: Client Access	4:3
Mailbox: Combined Hub/CAS	1:1

* <http://technet.microsoft.com/en-us/library/dd346701.aspx>

Table 45. Recommended Server Role Ratios Based on Processor Core*

© 2011 VMware, Inc. All rights reserved.

Exchange 2010 server role	Minimum Supported	Recommended
Hub Transport	4GB	1GB per core (4GB minimum)
Client Access	4GB	2GB per core (8GB minimum)
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)

* <http://technet.microsoft.com/en-us/library/dd346700.aspx>

Example:

Given our previous example of 16,000 mailboxes with 4,000 active mailboxes per server:

Hub Transport Calculations

Given that that we're running in a DAG, we need to think in terms of peak utilization of servers and the number of cores required to run ACTIVE mailboxes after a single node failure:

- 24 mailbox processor cores * .71 peak utilization = *18 mailbox processor cores* (rounded up).
- 18 mailbox processor cores/ 5:1 hub core ratio with AV = *4 Hub Transport processor cores* (rounded for even distribution from 3.6).

Based on your requirements, you may decide to deploy *2 Hub Transport virtual machines at 2 vCPU per virtual machine*.

To calculate memory for each virtual machine, multiply the number of vCPU * the recommended memory per core from the above table: 2 processor cores x 1GB per core = 2GB RAM; however, we must allocate *4GB of RAM* based on the minimum supported configuration.

Client Access Server Calculations

18 mailbox processor cores/4:3 Client Access core ratio = *16 Client Access processor cores* (rounded for even distribution from 14)

Based on your requirements, you may decide to deploy *4 Client Access virtual machines at 4 vCPU per virtual machine*.

To calculate memory for each virtual machine, multiply the number of vCPU * the recommended memory per core from the above table: 4 processor cores x 2GB per core = *8GB RAM*.

5.2.8 Resource Requirements by Server Role

For any application, first plan its resource requirements, then determine the underlying infrastructure requirements to meet the applications needs.

Table 46. Example Exchange Server Role Resource Requirements

Exchange Role	Physical Resources (per server)
Mailbox Server (4 servers)	CPU: 8 cores Memory: 64GB OS and Application File Storage: 80GB (OS & application files) Database and Log Storage 138 x 300GB/10K RPM FC/SCSI/SAS 3.5" Restore LUN Storage 9 x 300GB/10K RPM FC/SCSI/SAS 3.5" Network: 1Gbps
Client Access Server (4 servers)	CPU: 4 cores Memory: 8GB Storage: 24GB (OS & application files) Network: 1Gbps
Hub Transport Server (3 servers)	CPU: 2 cores Memory: 4GB Storage: 20GB (OS, application, & log files) 32GB (DB, protocol/tracking logs, & temp files) Network: 1Gbps

5.2.9 Guest Virtual Machine Configuration

The next step is to plan the individual virtual machine configuration based on the required resources for each Exchange server role.

Table 47. Exchange Virtual Machine Configuration

Exchange Role	Virtual Hardware (per VM)
Mailbox Server (4 VMs)	<p>CPU: 8 vCPU</p> <p>Memory: 64GB</p> <p>Storage: SCSI Controller 0</p> <p>HDD 1: 80GB (OS & application files)</p> <p>Storage: SCSI Controller 1</p> <p>HDD 2: 1321GB (DB1)</p> <p>HDD 3: 1321GB (DB2)</p> <p>HDD 4: 1321GB (DB3)</p> <p>HDD 5: 1321GB (DB4)</p> <p>HDD 6: 1321GB (DB5)</p> <p>HDD 7: 1321GB (DB6)</p> <p>HDD 8: 1321GB (DB7)</p> <p>HDD 9: 1321GB (DB8)</p> <p>HDD 10: 1321GB (DB9)</p> <p>HDD 11: 1321GB (DB10)</p> <p>HDD 12: 1321GB (DB11)</p> <p>HDD 13: 1321GB (DB12)</p> <p>Storage: SCSI Controller 2</p> <p>HDD 14: 1321GB (DB13)</p> <p>HDD 15: 1321GB (DB14)</p> <p>HDD 16: 1321GB (DB15)</p> <p>HDD 17: 1321GB (DB16)</p> <p>HDD 18: 1321GB (DB17)</p> <p>HDD 19: 1321GB (DB18)</p> <p>HDD 20: 1321GB (DB19)</p> <p>HDD 21: 1321GB (DB20)</p> <p>HDD 22: 1321GB (DB21)</p> <p>HDD 23: 1321GB (DB22)</p> <p>HDD 24: 1321GB (DB23)</p> <p>HDD 25: 1321GB (DB24)</p> <p>Storage: SCSI Controller 3</p> <p>HDD 26: 1206GB (Restore LUN)</p> <p>Network: NIC 1</p>

Client Access Server (4 VMs)	CPU: 4 vCPU Memory: 8GB Storage: SCSI Controller 0 HDD 1: 24GB (OS & application files) Network: NIC 1
Hub Transport Server (2 VMs)	CPU: 2 vCPU Memory: 4GB Storage: SCSI Controller 0 HDD 1: 20GB (OS and application files) Storage: SCSI Controller 0 HDD 2: 32GB (DB, protocol/tracking logs, temp files) Network: NIC 1

5.2.10 Planning the Host Hardware Configuration

5.2.10.1. Virtual Machine Distribution

In this example, we use four ESXi hosts connected to shared storage to use advanced VMware features such as HA and DRS.

To build infrastructure availability into the architecture, we distribute the 10 total virtual machines across four physical VMware ESXi hosts. Initial placement of virtual machines is relatively unimportant, especially if you're utilizing DRS.

Table 48. Exchange Virtual Machine Distribution

ESXi Host	Virtual Machines
ESXi Host 1	Exchange Mailbox VM 1 (8 vCPU/64GB RAM) Exchange Client Access VM 1 (4 vCPU/8GB RAM) Exchange Hub Transport VM 1 (2 vCPU/4GB RAM)
ESXi Host 2	Exchange Mailbox VM 2 (8 vCPU/64GB RAM) Exchange Client Access VM 2 (4 vCPU/8GB RAM) Exchange Hub Transport VM 2 (2 vCPU/4GB RAM)
ESXi Host 3	Exchange Mailbox VM 3 (8 vCPU/64GB RAM) Exchange Client Access VM 3 (4 vCPU/8GB RAM)
ESXi Host 4	Exchange Mailbox VM 4 (8 vCPU/64GB RAM) Exchange Client Access VM 4 (4 vCPU/8GB RAM)

5.2.10.2. ESXi Host Specifications

Each ESXi host should provide enough physical hardware resources to accommodate the planned workload and provide some headroom in the event of a HA failover or planned vMotion migration of live virtual machines for host hardware maintenance. The following table summarizes the ESXi host hardware configuration based on our example architecture.

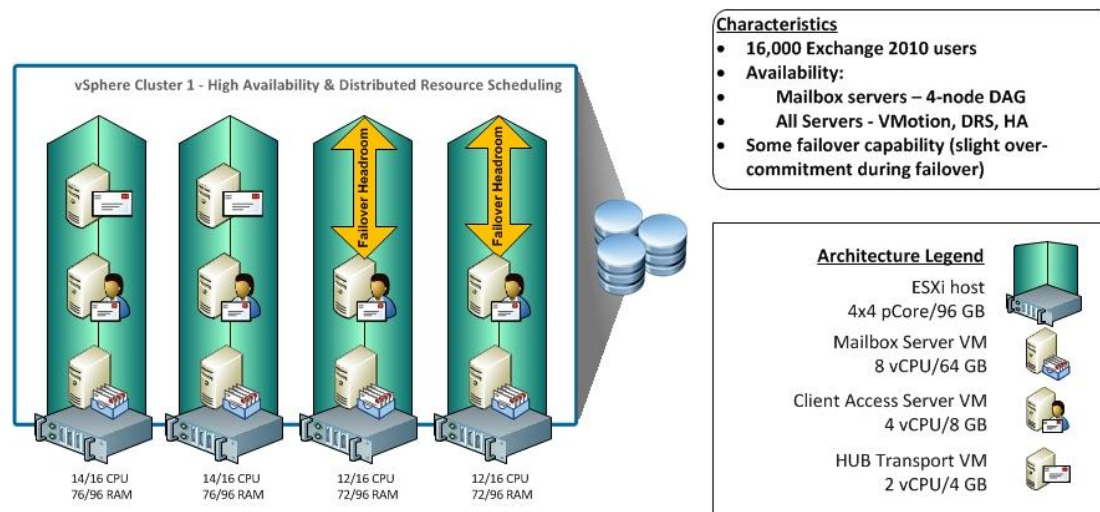
Table 49. Example ESXi Host Hardware Configuration Table

ESXi Host	Virtual Machines
ALL ESXi hosts	16 cores (4x4)
	96GB RAM (extra 20GB above requirements for use in failover)
	2 Fibre Channel HBAs
	4 Gigabit network adapters

5.2.10.3. Initial Virtual Machine Placement

Although the workloads migrate automatically with DRS (including the mailbox servers), Figure 10 is a useful planning tool for initial placement of virtual machines and calculating host failover capacity. At initial placement, ESXi hosts 3 and 4 have the most failover headroom.

Figure 10. Initial Virtual Machine Placement



6. vSphere Enhancements for Deployment and Operations

You can leverage vSphere to provide significant benefits in a virtualized Exchange datacenter, including:

- Increased operational flexibility and efficiency – Rapid software applications and services deployment in shorter time frames.
- Efficient change management – Increased productivity when testing the latest Exchange software patches and upgrades.
- Minimized risk and enhanced IT service levels – Zero-downtime maintenance capabilities, rapid recovery times for high availability, and streamlined disaster recovery across the datacenter.

6.1 VMware vMotion, DRS, and VMware HA

VMware vSphere vMotion technology enables the migration of virtual machines from one physical server to another without service interruption this migration allows you to move Exchange virtual machines from a heavily-loaded server to one that is lightly loaded, or to offload them to allow for hardware maintenance without any downtime.

VMware Distributed Resource Scheduler (DRS) takes the vMotion capability a step further by adding an intelligent scheduler. DRS allows you to set resource assignment policies that reflect business needs. DRS does the calculations and automatically handles the details of physical resource assignments. It dynamically monitors the workload of the running virtual machines and the resource utilization of the physical servers within a cluster.

VMware vMotion and DRS perform best under the following conditions:

- The source and target ESXi hosts must be connected to the same gigabit network and the same shared storage.
- A dedicated gigabit (or higher) network for VMware vMotion is recommended.
- The destination host must have enough resources.
- The virtual machine must not use physical devices such as CD ROM or floppy.
- The source and destination hosts must have compatible CPU models, or migration with VMware vMotion will fail. For a listing of servers with compatible CPUs, consult VMware vMotion compatibility guides from specific hardware vendors.
- To minimize network traffic it is best to keep virtual machines that communicate with each other together (for example, Exchange Mailbox and Global Catalog Servers) on the same host machine.
- Virtual machines with smaller memory sizes are better candidates for migration than larger ones.

Note As of Exchange 2010 SP1 Microsoft supports the use of VMware vMotion, and vSphere HA for Exchange 2010 DAG nodes. VMware has published best practices for combining these technologies in the whitepaper *Using HA, DRS, and vMotion with Exchange 2010 DAGs* (<http://www.vmware.com/files/pdf/solutions/VMware-Using-HA-DRS-vMotion-with-Exchange-2010-DAGs.pdf>).

With VMware High Availability (HA), Exchange virtual machines on a failed ESXi host can be restarted on another ESXi host. This feature provides a cost-effective failover alternative to expensive third-party clustering and replication solutions. If you use VMware HA be aware that:

- VMware HA handles ESXi host hardware failure and does not monitor the status of the Exchange services—these must be monitored separately.
- VMware HA heartbeat is sent via the vSphere VMkernel network, so redundancy in this network is recommended.

- Allowing two nodes from the same DAG to run on the same ESXi host for an extended period is not recommended when using symmetrical database distribution. DRS anti-affinity rules can be used to mitigate the risk of running active and passive databases on the same ESXi host.

6.2 Templates

VMware template cloning can increase productivity of system administration and testing in Exchange environments. A VMware template is a *golden image* of a virtual machine that can be used as a master copy to create and provision new virtual machines. It includes the guest operating system and Exchange application data. You can use virtual machine templates to provision a new preconfigured Exchange system. In native environments, this process can consume significant time, requiring you to procure hardware and install the operating system. Cloning produces a controlled virtual machine configuration so deployment is less error prone and less time-consuming.

6.3 VMware vCenter Operations

VMware vCenter Operations™ can provide a holistic approach to performance, capacity, and configuration management. By using patented analytics, service levels can be proactively monitored and maintained. When performance or capacity problems arise in your Exchange environment vCenter Operations is able to analyze metrics from the application all the way through to the infrastructure to provide insight into problematic components, whether they are compute (physical or virtual), storage, networking, OS, or application-related. By establishing trends over time vCenter Operations can cut through the noise of false alerts and proactively alert on the potential root cause of increasing performance problems before end users are impacted.

In an Exchange environment constant monitoring is required to maintain acceptable service levels, not only for end users but for the Exchange components as well. vCenter Operations includes patented capacity analytics that can eliminate the need for spreadsheets, simple scripts, or rules of thumb. Quickly run through “what if” capacity scenarios to understand growth trends and identify upcoming compute power shortages or over-provisioned resources. As an application comprised of multiple components Exchange performance and functionality can be affected by changes made to any of the roles. vCenter Operations monitors configurations across virtual machines and detects unwanted changes to help maintain continuous compliance with operational best practices.

Figure 11. VMware vCenter Operations



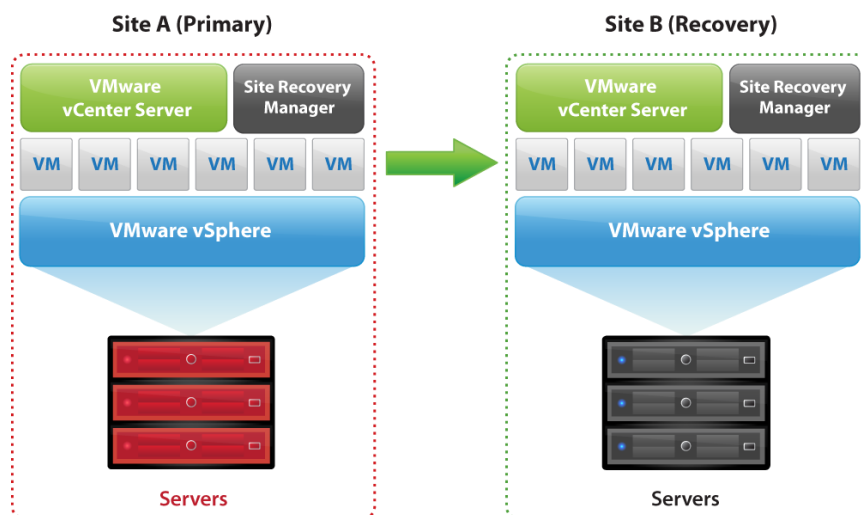
6.4 VMware vCenter Site Recovery Manager

VMware vCenter Site Recovery Manager (SRM) takes advantage of virtual machine encapsulation to make testing and initiating DR failover a straightforward, integrated vCenter process. SRM runs alongside VMware vCenter Server to provide planning, testing, and automated recovery in the case of a disaster. By using vSphere Replication or storage-based replication technology vCenter SRM eliminates the manual steps required during a failover scenario to provide consistent and predictable results each time. At a high level the steps that can be performed during a failover test or actual run are as follows:

- Shut down production virtual machines (failover).
- Promote recovery storage to primary (failover).
- Take and mount snapshot of recovery storage in read/write mode (test only).
- Rescan ESXi hosts to make storage visible.
- Register recovery virtual machines.
- Power-on virtual machines at recovery site.
- Reconfigure IP settings and update DNS if required.
- Verify VMware Tools starts successfully on recovered virtual machines.
- Power-off recovered virtual machines (test only).
- Un-register virtual machines (test only).
- Remove storage snapshot from recovery side (test only).

With Exchange 2010 comes a single method for providing protection from unplanned downtime, database availability groups (DAGs). A DAG can provide automated failover in the case of a database, guest, or even host failure. A DAG can also provide manually activated protection by implementing *Lagged Log Replay*, which is usually deployed to establish disaster recovery. DAGs require Windows Enterprise licenses and storage for each copy of a protected database. Though an excellent choice for datacenter high-availability, the application-centric nature of a DAG may not be in line with a company's disaster recovery plans. vCenter SRM is not a replacement for application-aware clustering solutions that may be deployed within the guest OS. SRM provides integration of the replication solution, VMware vSphere, and optionally customer-developed scripts to provide a simple, repeatable, and reportable process for disaster recovery of the entire virtual environment, regardless of the application.

Figure 12. VMware vCenter Site Recovery Manager



Case Study:

VMware IT protects Exchange using vCenter Site Recovery Manager.

Challenge

Create and implement a disaster recovery plan in conjunction with a datacenter consolidation project.

With the consolidation of two primary datacenters VMware IT was required to consolidate its virtualized Exchange environment from a site-resilient design into the foot-print available at a single datacenter. The challenge was how to maintain an additional online copy of the data that could quickly and easily be activated for disaster recovery purposes.

Solutions

- Implement array-based replication between the primary datacenter in Palo Alto, CA. and the DR datacenter in Wenatchee, WA.
- Deploy a vCenter Server and vCenter Site Recovery Manager in both locations.
- Create SRM recovery plans for each individual Exchange mailbox server and an all-encompassing recovery plan for all Exchange mailbox servers.

For its Exchange email environment VMware IT chose to deploy an array-based replication solution supported for use by vCenter Site Recovery Manager. With asynchronous replication between the two datacenters established, the IT team deployed parallel vCenter Server and SRM services running atop the production virtual infrastructure in each location. The goal was that this solution could be used not only for disaster recovery, but in the event of a single mailbox server failure as well. To accomplish this, SRM recovery plans were created for each individual Exchange mailbox virtual machine. In a true disaster scenario all Exchange mailbox virtual machines would have to be recovered. Although each recovery plan could be initiated consecutively, a single recovery plan incorporating all Exchange mailbox virtual machines was created. This allows the vSphere or Exchange administrator to initiate one recovery plan and focus on other tasks while the mailbox virtual machines were brought online at the DR facility.

Results

- Recovery Point Objective of less than 5 minutes.
- Recovery Time Objective of under 60 minutes.
- Automated manual recovery steps such as reconfiguring IP addresses, updating DNS, promoting and reversing replication of storage arrays.

VMware IT is today able to achieve a recovery point objective (RPO) of less than five minutes using the implemented asynchronous array replication solution. Using vCenter Site Recovery Manager to automate tasks such as promoting the secondary storage to primary, registering virtual machines, and reconfiguring IP settings among others, VMware IT is able to achieve a recovery time objective (RTO) of about 60 minutes.

More information

More information about the VMware implementation of SRM can be found in *VMware Professional Services Helps Architect Comprehensive Disaster Recovery Plan*
http://www.vmware.com/files/pdf/customers/VMW_10Q1_CS_PSO_DisasterRecoveryPlan_USLet_EN.pdf

6.5 VMware vCenter AppSpeed

VMware vCenter AppSpeed provides a real-time view of end-user application performance without impacting overall system performance. AppSpeed inspects traffic flowing over the vSwitch to discover and map the applications and protocols found traversing the virtual network. The data collected by AppSpeed is then used to measure actual performance against SLAs and enables root cause analysis. Because the AppSpeed server and probes are deployed as virtual appliances they integrate directly into your existing vCenter environment.

When virtualizing a multitier application such as Exchange 2010, knowing the virtual network trends between the hub transport, client access, and mailbox layers can assist in establishing baseline trends and SLAs. Each tier or role in an Exchange 2010 environment has dependencies with other Exchange roles, Active Directory, and DNS. Having the ability to map these dependencies can reduce the time needed to troubleshoot critical activities such as:

- MAPI mail submission between the hub transport and mailbox servers.
- Latency for Outlook Web App and Exchange Web Services.
- MAPI latency between the client access and mailbox server roles.
- RPC latency between the clients and the client access pool.

Figure 13. VMware vCenter AppSpeed

