

WHITE PAPER

Best Practices Guide: Microsoft Exchange Solutions on VMware



Contents

1. Introduction	3
1.1. Overview	3
1.2. Purpose.....	3
1.3. Target Audience	3
1.4. Scope	4
2. VMware Infrastructure Best Practices for Exchange	5
2.1. Physical and Virtual CPUs	5
2.2. Memory Configuration Guidelines	6
2.3. Storage Virtualization.....	9
2.4. Networking Best Practices.....	13
3. Exchange Server 2007 Performance	16
3.1. Overview	16
3.2. Key Performance Considerations	17
3.3. Performance Testing	17
3.4. Ongoing Performance Monitoring and Tuning	18
4. Sample Exchange Server 2007 Sizing and Configuration Process	20
4.1. Understanding Workload Requirements	20
4.2. Understanding Microsoft Sizing Guidelines	20
4.3. Exchange Mailbox Server Building Blocks	21
4.4. Sample Building Block Sizing – Average Profile for 4,000 Users	24
4.5. Sizing Peripheral Server Roles	29
4.6. Planning Guest Virtual Machine Configurations	30
4.7. Planning the Host Hardware Configuration	31
4.8. VMware ESX Host Architecture	33
4.9. Logical Architecture.....	34
5. Deployment and Operations.....	35
5.1. VMware Infrastructure Deployment Features.....	35
5.2. VMware VMotion, VMware DRS, and VMware HA	36

1. Introduction

1.1. Overview

E-mail 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 2007, Microsoft has added many features that improve messaging performance, reliability, and scalability. These provide a major step forward. However, Exchange Server 2007 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 2007 and the enhancements enabling it to run on the 64-bit Windows architecture 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® Infrastructure can accomplish these goals.

1.2. Purpose

This whitepaper will provide best practice guidelines for deploying Exchange Server 2007 on VMware Infrastructure. The recommendations in this guide are not specific to any particular set of hardware or to the size and scope of the Exchange implementation. The examples and considerations in this document provide guidance only and do not represent strict design requirements, as the flexibility of Exchange Server 2007 on VMware Infrastructure allows for a wide variety of valid configurations.

1.3. Target Audience

This guide assumes a basic knowledge and understanding of VMware Infrastructure and Exchange Server 2007.

Architectural staff can reference 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 reference 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.

1.4. Scope

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

- **VMware Infrastructure Best Practices for Exchange** – This section provides best practice guidelines for ensuring that the VMware Infrastructure platform is properly prepared for running Exchange Server 2007. This section includes guidance in the areas of CPU, memory, storage, and networking.
- **Exchange Server 2007 Performance** – This section provides background information on Exchange Server 2007 performance in a virtual machine. It also provides information on official VMware partner testing and guidelines for conducting and measuring internal performance tests.
- **Sample Exchange Server 2007 Sizing and Configuration Process** – Sizing Exchange to run in a virtual machine follows many of the same best practices as sizing on physical servers. This section walks the reader through sizing a sample configuration, starting with the mailbox server virtual machine and adding the peripheral server roles based on Microsoft guidelines.
- **Deployment and Operations** – This section provides best practice guidelines for deploying an Exchange Server 2007 environment on VMware Infrastructure, as well as a brief look at operational enhancements made available by VMware.

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

- **Sizing and Configuration Samples** – This information can be found in the **Design and Sizing Examples** document included in this Solution Kit.
- **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 the **Availability and Recovery Options Solution Brief** included in this Solution Kit.
- **Tactical Steps for Deployment** – Detailed deployment steps, including template-based deployment of Exchange Server 2007 virtual machines, can be found in the **Deployment Guide** 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 VMware Infrastructure. Exchange deployments cover a wide subject area, and Exchange-specific design principles should **always** follow Microsoft guidelines for best results.

2. VMware Infrastructure Best Practices for Exchange

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

2.1. 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-based processors. 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 begins executing instructions from the virtual machine.

2.1.1. Virtual SMP

VMware Virtual SMP enhances virtual machine performance by enabling a single virtual machine to use multiple physical processors simultaneously. 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 benefit from SMP.

The virtual processors from SMP-enabled virtual machines are co-scheduled. That is, if physical processors 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. The following points should be considered when using multiple vCPUs:

- 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, VMware ESX Server manages these “idle” processors, resulting in some additional work by the hypervisor. These low-utilization vCPUs compete with other vCPUs for system resources.

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 (i.e. production systems), try to ensure the total number of vCPUs assigned to all the virtual machines is equal to or less than the total number of cores on the VMware ESX Server host machine.

2.1.2. Service Console

The VMware ESX Server service console provides an execution environment to monitor and administer the entire VMware ESX Server host. The service console operating system (COS) is a reduced version of Red Hat Enterprise Linux. The COS is scheduled on CPU 0 for VMware ESX Server version 3.0.x. Under high loads the COS will compete with virtual machines that are scheduled on CPU 0. VMware recommends you avoid running programs in the COS. The elimination of COS in VMware ESX Server 3i removes this contention.

2.2. Memory Configuration Guidelines

This section provides guidelines for configuring the memory requirements of multiple Exchange virtual machines on a single server to maximize performance. The guidelines outlined here take into account VMware ESX Server memory overhead and the virtual machine memory settings.

VMware ESX Server 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 (i.e. same guest operating systems);
- Memory over-commitment, which means allocating more memory to virtual machines than is physically available on the VMware ESX Server 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 on VMware ESX Server memory management concepts, consult the VMware Resource Management Guide (http://www.vmware.com/pdf/vi3_301_201_resource_mgmt.pdf).

2.2.1. Memory Overhead

VMware ESX Server requires three areas of memory overhead:

- **A fixed, system-wide overhead for the service console** (about 272 MB for VMware ESX 3.x).
- **A fixed, system-wide overhead for the VMkernel**, which depends on the number and size of device drivers.
- **Additional overhead for each virtual machine**, since the virtual machine monitor for each virtual machine requires some memory for its code and data.

You can view the fixed, system-wide overhead for any VMware ESX Server instance in the VMware Infrastructure Client. Select the **Configuration** tab under **Memory Settings**.

The overhead for each virtual machine, noted in the table below, depends on the number of vCPUs, the amount of memory allocated to the virtual machine, and whether the virtual machine is 32-bit or 64-bit.

Virtual Machine Memory Overhead¹

Virtual CPUs	Memory (MB)	Overhead for 32-bit virtual machine (MB)	Overhead for 64-bit virtual machine (MB)
1	256	79	174
1	512	79	176
1	1024	84	180
1	2048	91	188
1	4096	107	204
1	8192	139	236
1	16384	203	300
2	256	97	288
2	512	101	292
2	1024	101	300
2	2048	125	316
2	4096	157	349
2	8192	221	413
2	16384	349	541
4	256	129	511
4	512	133	515
4	1024	141	523
4	2048	157	540
4	4096	189	572
4	8192	222	605
4	16384	350	734

2.2.2. Virtual Machine Memory Settings


The VMware ESX Server memory settings for a virtual machine include the following parameters:

- The configured size, which is the memory size specified at creation of the virtual machine.
- The reservation, which is a guaranteed lower bound on the amount of memory that the host reserves for the virtual machine, and cannot be reclaimed by VMware ESX Server for other virtual machines.

¹ (Source: *VMware Resource Management Guide*. Note second column refers to memory allocated to VM.)

The figure below illustrates the memory settings for a virtual machine:

Virtual Machine Memory Settings



- **Configured memory** = memory size of virtual machine assigned at creation.
- **Touched memory** = memory actually used by the virtual machine. VMware ESX Server only allocates guest memory on demand.
- **Swappable** = virtual machine memory that can be reclaimed by the balloon driver or by VMware ESX Server swapping. Ballooning will occur before VMware ESX Server swapping. If this memory is in use by the virtual machine (i.e. 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, VMware ESX Server will forcibly reclaim memory from the virtual machine using the **VMkernel swap file**.

2.2.3. Sizing Memory of Multiple Exchange Virtual Machines

Based on the memory overhead defined above, the available memory for Exchange applications (and the guest operating systems) in virtual machines is equal to the total VMware ESX Server host physical memory minus system-wide memory overhead, virtual machine memory overhead, and a user-defined "memory buffer."

Note: The memory buffer is not a VMware parameter but a user-defined value designed to provide headroom and flexibility to manage more virtual machines than the initial estimate (i.e. virtual machines migrated using VMware VMotion onto the system from another VMware ESX Server host machine). This will depend on customer-specific design requirements.

VMware ESX Server page sharing makes more physical memory available, but this additional memory is not counted here to provide a more conservative estimate. As Exchange servers are memory intensive and performance is often a key factor (e.g. in production environments), VMware recommends the following practices:

- Do not over-commit memory on the VMware ESX Server host until VMware VirtualCenter reports that steady-state memory usage is below the amount of memory on the server.
- Set the memory reservation to the configured size of the virtual machine, which results in a per-virtual machine VMkernel swap file of zero bytes. The guest operating system within the virtual machine will maintain its own separate swap/page file.
- Do not disable the balloon driver (which is installed with VMware Tools).

- Size the memory of each virtual machine so that the total configured memory for all virtual machines is equal to or less than the memory available for Exchange applications as calculated above.
- 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 will be invoked as in normal, native operations. Memory and swap/page file configuration for Exchange virtual machines follows the same guidelines as for native environments. In general, these should be set to minimize any guest OS swapping.
- Microsoft documentation and recommendations should be closely followed; key points will be identified in the following sections.

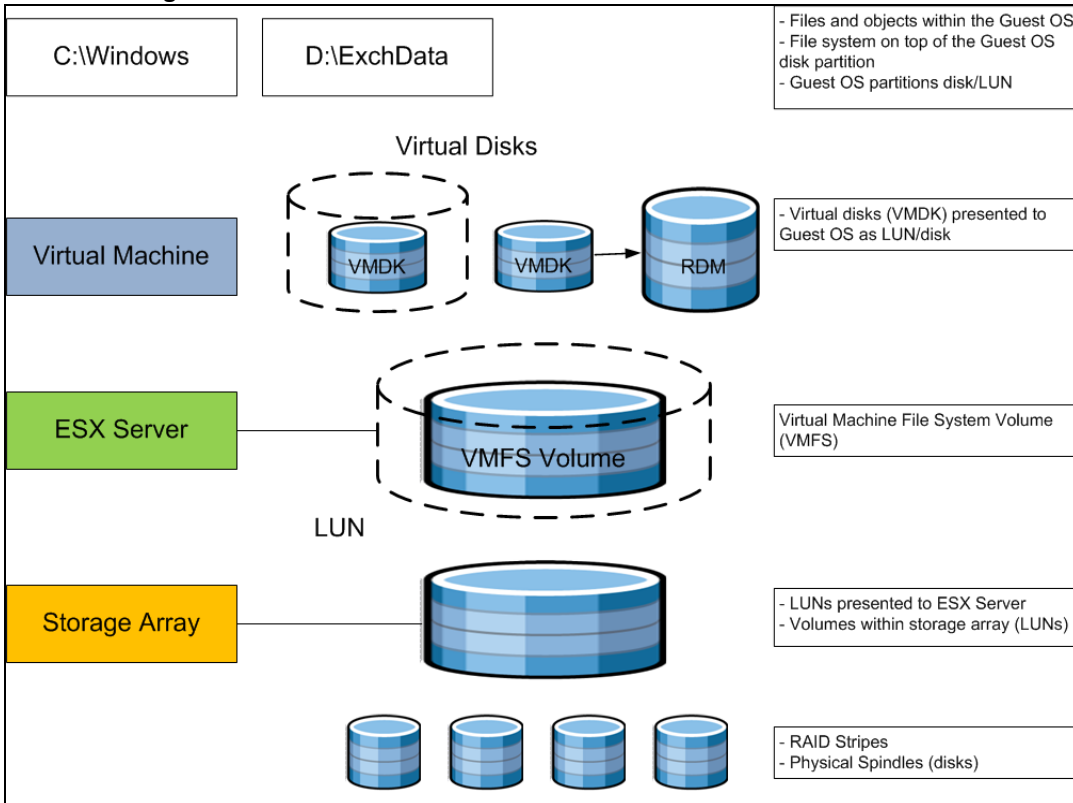
2.3. Storage Virtualization

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 machines files on shared storage to take advantage of VMware 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 the figure below. 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 VMware ESX Server. Storage array LUNs 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.


VMware Storage Virtualization Stack



2.3.1. Storage Multipathing

VMware recommends you set up a minimum of four paths from a VMware ESX Server host to a storage array, which means the host requires at least two HBA ports.

Storage Multipathing Requirements for VMware ESX Server



The terms used in the figure above are defined here:

- **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, please consult the [VMware SAN System Design and Deployment Guide](#).

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. • Fully supports VMware Site Recovery Manager. 	<ul style="list-style-type: none"> • Maps a single LUN to one virtual machine so only one virtual machine is possible per LUN. • More LUNs are required, so it is easier to reach the LUN limit of 256 that can be presented to ESX Server. • Uses RDM to leverage array-level backup and replication tools integrated with Exchange databases. • Although not required, RDM volumes can help facilitate moving Exchange data from virtual to standby physical boxes in certain support circumstances. • Required for third-party clustering software (e.g. MSCS). Cluster data and quorum disks should be configured with RDM. • Some customers use RDMS for Exchange databases and logs on the MBX server role to guarantee that no other VMs are provisioned to those LUNs. • Experimental support for VMware Site Recovery Manager.

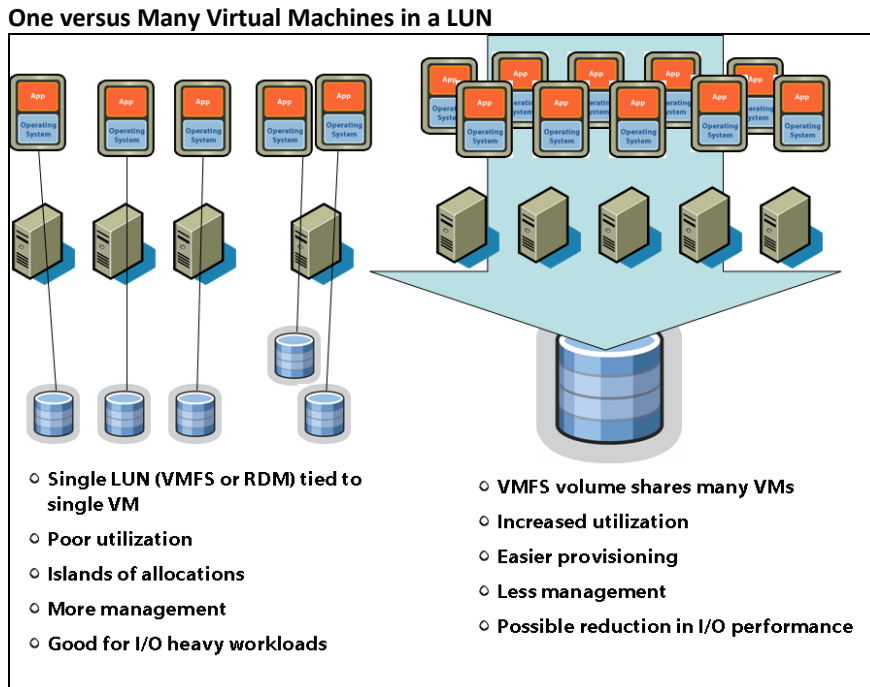
It is also possible and even advantageous in some circumstances to mix VMFS and RDM in Exchange environments under the following conditions:

- 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.
- 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.

It is important to note that there are several different shared-storage options available (iSCSI, Fiber Channel, NAS, etc.). To see the most recent list of compatibilities please consult the latest [VMware Compatibility Guides](#).

2.3.3. Number of Virtual Machines per LUN

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



2.4. Networking Best Practices

This section covers design guidelines for the virtual networking environment and provides configuration examples at the VMware ESX Server host level for Exchange Server 2007 installations.

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

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, VMware ESX Server 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 VMware Infrastructure 3. 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. Virtual Networking Concepts

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 VMware ESX Server 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 VMware ESX Server enables port-based VLAN tagging based on the switch ports. The [VMware Virtual Networking Concepts](#) document (available for download at 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 VMware ESX Server 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 VMware ESX Server are templates for creating virtual ports with a particular set of specifications. In VMware ESX Server, there are three types of port group/virtual switch connections:

- Service console port group: VMware ESX Server management interface.
- VMkernel port group: VMware VMotion, iSCSI, and/or NFS/NAS networks.
- Virtual machine port group: virtual machine networks.

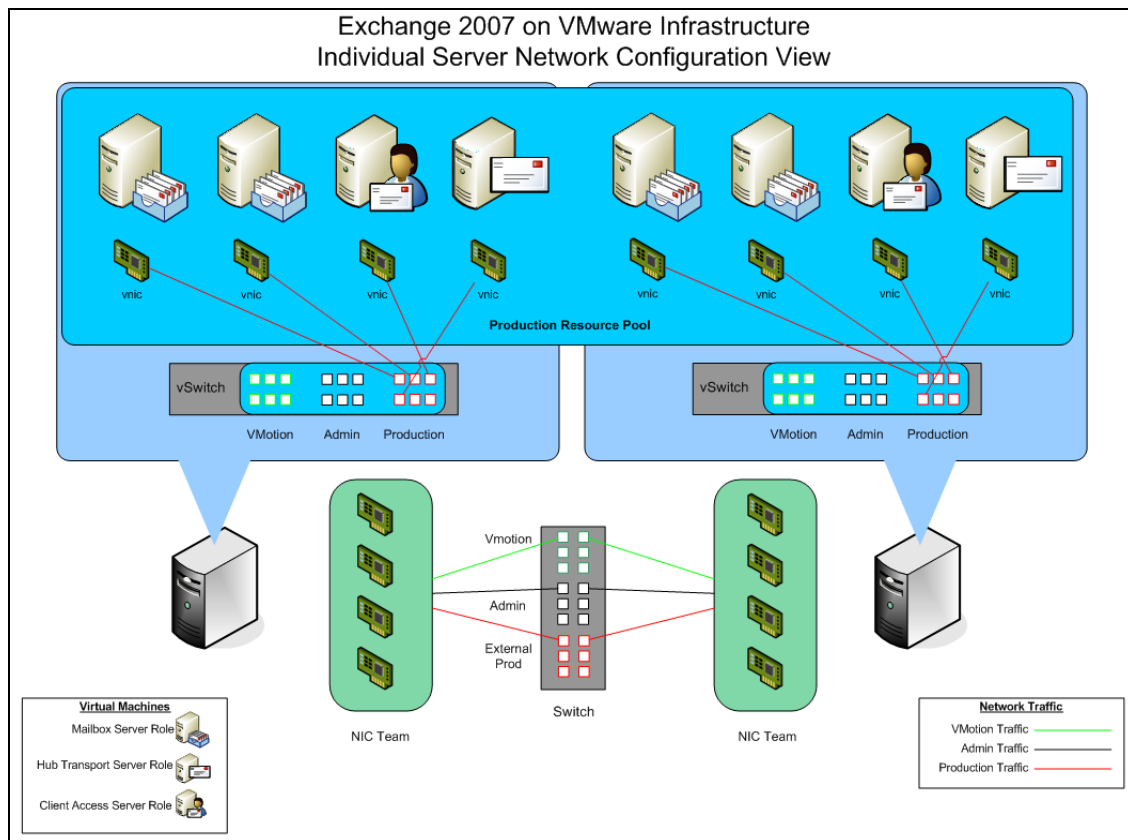
More than one connection type can exist on a single virtual switch, or each connection type can exist on its own virtual switch.

2.4.1.2. NIC Teaming

A single virtual switch can be connected to multiple, physical Ethernet adapters using the VMware Infrastructure feature called NIC teaming. This provides redundancy and/or aggregation.

2.4.2. Sample Exchange Virtual Network Configuration

The following figure is an example of what a network layout for an Exchange production environment might look like:



This figure illustrates how networking is handled at the VMware ESX level. Each VMware ESX Server 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, you will want to make sure that production environment network traffic remains separate from VMware 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 VMware ESX server. This network traffic is handled through the virtual switch.

Each physical NIC can be configured to connect directly to an assigned VLAN, but the VMware VMotion and Admin networks are not used heavily as production networks. One practice is to team all the NICs on the VMware ESX server, 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 VMware VMotion and Admin VLANs are not in heavy use.

3. Exchange Server 2007 Performance

3.1. Overview

VMware and its partners have used testing to successfully demonstrate the viability of running Exchange 2003 on VMware Infrastructure 3. This testing has been confirmed by organizations that have deployed Exchange 2003 in virtualized production environments and now benefit from its operational advantages. Some have virtualized their entire Exchange 2003 environment and have carefully designed their infrastructure to accommodate the architectural challenges previously noted. Others have split their Exchange implementations, limiting virtualization to server roles such as bridgehead servers, SMTP gateways, and infrastructure services such as domain controllers and global catalogue servers.

Exchange Server 2007, however, is proving to be much easier to deploy and operate than its predecessors. The dynamics of deploying Exchange Server 2007 are shifting dramatically in ways that enhance the benefits offered by the VMware Infrastructure 3 platform. This shift is a result of advancements in three key areas:

1. **Architectural improvements of Exchange Server 2007 on a 64-bit platform** 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. With full support for 64-bit Windows virtual machines, these same benefits are realized when deploying Exchange Server 2007 on VMware Infrastructure 3.
2. **Advances in server hardware** such as dual and quad-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. Recent performance studies have shown unprecedented performance and scalability of up to 16,000 users on a single 64-bit physical server when Exchange Server 2007 is deployed on VMware Infrastructure 3.
3. The advances in Exchange Server 2007 and server hardware technology have coincided with advances in **VMware Infrastructure 3**. Virtual machines can now support up to 64 GB RAM and four vCPUs, and are capable of running even the largest Exchange mailbox servers. VMware ESX servers can now take advantage of up to 256 GB of physical RAM. Network improvements such as Jumbo Frames and TCP Segment Offload have lowered overall CPU usage. These and other enhancements make VMware Infrastructure 3 capable of meeting performance requirements for even the most demanding Exchange workloads.

Third-party testing of Exchange Server 2007 in virtual operation has been completed with Microsoft's JetStress and LoadGen tools, the standard tools for Exchange performance analyses. 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 2007 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 2007 performance on VMware Infrastructure 3, including processor and memory allocation to the guest virtual machine, storage layout/design, virtual machine placement, and application clustering methods, to name a few. Below are some tips for ensuring 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](#).
- 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 of this document to ensure that the VMware Infrastructure environment is optimized for enterprise applications such as Exchange.

3.3. Performance Testing

Every Exchange environment is different, with varying business and technical requirements, a plethora of server and storage options, and requirements for integrating with third-party software solutions such as anti-virus, anti-spam, PDA messaging, etc. Due to the many variables, it is highly 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 VMware Infrastructure. Both of these options will be discussed in this section.

3.3.1. Internal Performance Testing

Microsoft provides tools to measure the performance of Microsoft Exchange Server architectures. LoadSim is used to measure Exchange Server 2003 performance and LoadGen is used to measure Exchange Server 2007 performance. 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 LoadSim and 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, and has been known to display a positive bias of 10-20 percent. This means that JetStress often reports 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.

3.3.2. Partner Performance Testing

VMware and its OEM partners have conducted several performance studies to characterize how the Exchange Server 2007 Mailbox Server role performs as a virtual machine on a variety of hardware platforms. Below are the results of those tests:

Partner	Type	Summary	Resource
Dell (server) & EqualLogic (storage)	Storage Sizing & Performance	Whitepaper compares physical versus virtual Exchange 2007 performance on Dell servers and EqualLogic storage.	http://www.vmware.com/files/pdf/exchange_equallogic_performance_wp.pdf
Dell (server) & EMC (storage)	Performance	Whitepaper compares physical versus virtual Exchange 2003 performance on Dell servers and EMC storage.	http://www.vmware.com/files/pdf/exchange_ms_virtualizing_wp.pdf
Dell (server)	Performance	Whitepaper describes performance and scalability of Exchange Server 2007 on the VMware Infrastructure platform.	http://www.vmware.com/files/pdf/exchange_dell_virtualizing.pdf
HP (server & storage)	Server/Storage Sizing & Performance	Whitepaper compares physical versus virtual Exchange Server 2007 performance on HP servers and storage.	http://www.vmware.com/files/pdf/exchange_hp_performance_wp.pdf

3.4. Ongoing Performance Monitoring and Tuning

Traditional Exchange Server performance monitoring leverages Microsoft's Windows performance monitor tool PerfMon to collect statistics. Exchange integrates with PerfMon to provide familiar counters that indicate system performance. However, 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 VMware ESX Server.

Exchange administrators should pay close attention to the following counters, referring to [online documentation on performance monitoring and analysis](#) for more information on these counters and their interpretation.

Performance Counters of Interest to Exchange Administrators

Subsystem	esxtop Counters	VirtualCenter Counter
CPU	%RDY %USED	Ready Usage
Memory	%ACTV SWW/s SWR/s	Active Swapin Swapout
Storage	ACTV DAVG/cmd KAVG/cmd	Commands deviceWriteLatency & deviceReadLatency kernelWriteLatency & kernelReadLatency
Network	MbRX/s MbTX/s	packetsRx 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 mis-configured 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.

4. Sample Exchange Server 2007 Sizing and Configuration Process

4.1. Understanding Workload Requirements

A fundamental component to successfully configuring an Exchange Server 2007 infrastructure is to understand the workload requirements. The key to quantifying this workload is to break down resource demand generation in Exchange to its most atomic level, the end user. Microsoft provides guidance on categorizing end user demand generation in “knowledge worker” profiles, or usage profiles. The following table shows the usage profiles used to plan Exchange Server 2007 infrastructure requirements.

User Profiles for Microsoft Outlook

User type (usage profile)	Sent/received per day (approximate 50-kilobyte (KB) message size)
Light	5 sent/20 received
Average	10 sent/40 received
Heavy	20 sent/80 received
Very heavy	30 sent/120 received

[http://technet.microsoft.com/en-us/library/aa998874\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/aa998874(EXCHG.80).aspx)

Based on the table above there are three key metrics necessary to classify your organization’s usage profile:

- Messages sent per user per day;
- Messages received per user per day;
- Average message size.

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 2007 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](#). 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. Understanding Microsoft Sizing Guidelines

The Microsoft sizing guidelines are broken down into four major categories for the purpose of infrastructure capacity planning:

- Processor Configuration
- Memory Configuration
- Storage Configuration
- Server Roles Ratios

These guidelines are documented on Microsoft TechNet in the “Planning Your Server and Storage Architecture” section of the Exchange Server 2007 documentation. VMware recommends that you follow Microsoft’s guidelines combined with your organizational standards to meet both business and functional requirements.

Each Exchange Server 2007 server role (Mailbox, Client Access, Hub Transport, Edge Transport, and Unified Messaging) has specific guidelines under each one of these categories for appropriately planning the necessary resources. There is a relationship between each resource component and the various server roles. We have broken down this relationship into a simple workflow to follow when planning your Exchange Server 2007 resource requirements on a VMware Infrastructure. These steps are identical to planning the Exchange Server 2007 resource requirements on a physical server infrastructure.

4.3. Exchange Mailbox Server Building Blocks

The building block approach is a recommended best practice for creating a virtualized Exchange environment running on VMware Infrastructure 3 using pre-sized virtual machine configurations. Exchange servers that have been broken down 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.

VMware and its partners have performed testing that focuses on four primary sizes for mailbox virtual machine building blocks consisting of 500, 1000, 2000, and 4000 users. Testing has validated performance and sizing for these configurations, and each has known performance profiles that can be leveraged for easy server sizing as well as for easy scalability.

The table below presents some guidelines for sizing the mailbox server role in a virtual machine using the building block approach for mailboxes with “average” user IOPS profiles. The same principles are used for sizing the other profiles.


Building block CPU and RAM Sizing for Mailboxes with Average Profiles

Building Block	500	1000	2000	4000
Profile	Average	Average	Average	Average
vCPU	1	1	2	4
Base RAM (GB)	2	2	2	2
Variable (MB/Mailbox)	<u>3.5MB/Mailbox</u>	<u>3.5MB/Mailbox</u>	<u>3.5MB/Mailbox</u>	<u>3.5MB/Mailbox</u>
Total (GB)	4	6	9	16

<http://www.microsoft.com/technet/prodtechnol/exchange/2007/plan/hardware.mspx>

For example: Company XYZ is going to design their mailbox server environment on VMware Infrastructure 3 to accommodate the 16,000 mailboxes in their organization. The IT team has decided to split the mailboxes into four separate virtual machines. However, all four virtual machines will run on the same ESX server, which contains 16 cores and 64 GB RAM. **The figure below** illustrates how four mailbox servers can be deployed on a single server:

16,000-user Building Block Configuration with Four Mailbox Virtual Machines



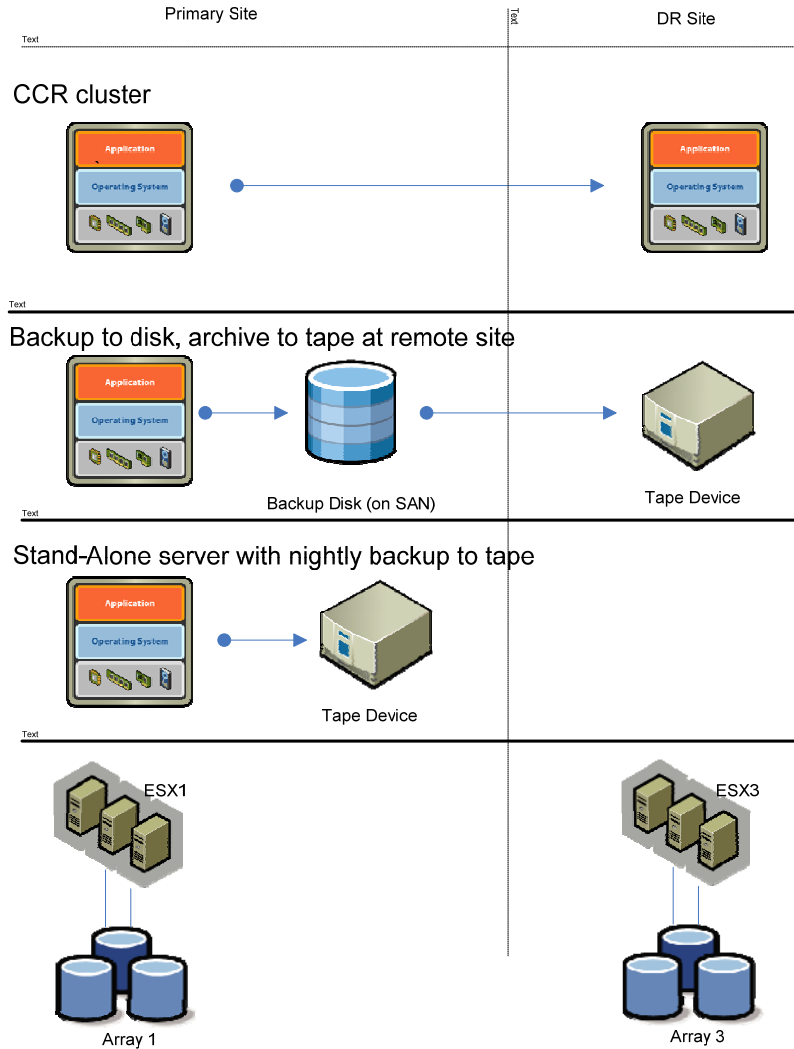
The building block approach offers the following benefits:

- **Reduced server hardware expense:** XYZ Corp can separate their users into four mailbox server virtual machines for maximum flexibility but will only require a single server. Without virtualization, this kind of configuration would require four separate servers.
- **Better use of powerful 64-bit servers:** Today's servers ship with multi-core processors and increased memory density. Running multiple Exchange virtual machines is one of the best ways to avoid under-utilization of CPU and RAM resources and unlock the full value of these systems.
- **Reduced risk:** An outage on a single mailbox virtual machine will affect fewer users since the users are spread across four separate mailbox servers.
- **Improved scalability:** Adding another mailbox server to the environment can be achieved in minutes. Match up the required number of users with the appropriate building block, deploy a new virtual machine from the template, and install and configure the appropriate Exchange server role.
- **Improved flexibility:** Each mailbox server virtual machine can be moved onto other VMware ESX servers at any time with no loss of service using VMware VMotion. Exchange servers are no longer tied to the underlying physical server, which creates a flexible Exchange environment that can easily change to meet changing business requirements.
- **Memory page sharing:** Using building blocks for Exchange presents some benefits in the way physical memory is used in the form of page sharing. Running multiple virtual machines on a VMware ESX Server with the same OS and Exchange binaries will benefit from page sharing memory with the VMware ESX Server architecture. This allows for more efficient use of the physical memory on a server.

- Match building blocks with specific Service Level Agreements (SLAs):** Having users spread across multiple building blocks allows an organization to meet different SLAs for different groups of users. In our example with 16,000 users across four mailbox virtual machines (supporting 4,000 users each), each virtual machine can be matched with its own backup, high availability, disaster recovery, and archive or other solutions to meet SLAs.

The figure below shows three mailbox virtual machines on a single VMware ESX server, each with its own set of requirements for meeting SLAs:

Using Virtual Machine Building Blocks to Meet Varying SLAs



When using the building block approach for sizing virtual machines, note the following:

- Smaller virtual machines (CPU and RAM) can be moved faster with VMware VMotion than larger virtual machines. For example, a mailbox server virtual machine with 2 vCPU and 9GB RAM can move to another VMware ESX Server with VMware 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 VMware Infrastructure 3 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 to a virtual machine 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.

4.4. Sample Building Block Sizing – Average Profile for 4,000 Users

Using the Microsoft sizing guidelines and the building blocks approach, we will apply the formula to size a 16,000-user environment with average mailbox profiles. The mailboxes will be 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 and are not intended for use in production deployments.

In our example, we will be using the average mailbox profile as defined below:

- 10 messages sent/40 received per day;
- Average message size of 50 KB;
- 500 MB mailbox quota.

Every environment is different and some organizations use email more heavily than others. To accurately determine your mailbox profile requirements, utilize the [Microsoft Exchange Server Profile Analyzer](#). It is also highly recommended that you work with a partner that is extremely well versed in Exchange Architecture to ensure proper performance in your specific environment.

Note that this example does not take into account a particular storage solution. Many of VMware's storage partners have performed extensive testing on building blocks of varying capacities and workloads. Please refer to the Partner Resource Catalog for storage-specific implementation details.

4.4.1. Building Block CPU Requirements (4,000 Users)

Based on your organization's usage profile, Microsoft provides guidance to estimate the number of processor cores to support your organization's Exchange workload. The current guidelines are:

- 1000 average profile users per processor core;
- 500 heavy profile users per processor core;
- Up to eight processor cores maximum.

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

Example:

Your organization plans to support 16,000 users with 4,000 mailboxes per Exchange Server 2007 mailbox server and a 500 MB mailbox quota. You have used the Exchange Server Profile Analyzer to determine that your users average 10 messages sent and 40 messages received per day, with an average message size of 50 KB. This translates to an average usage profile.

- 16,000 mailboxes/4,000 mailboxes per server = **4 mailbox servers**
- 4000 mailboxes per server/1,000 average users per core = **4 processor cores per server**
- 4 mailbox servers x 4 processor cores per server = **16 processor cores**

4.4.2. Building Block Memory Requirements (4,000 Users)

Based on your organization's usage profile, Microsoft provides guidance to estimate the amount of RAM to support your organization's Exchange workload. The current guidelines are shown in the following table:

Mailbox Server Memory Recommendations

User type	Mailbox server memory recommendation
Light	2 GB plus 2 MB per mailbox
Average	2 GB plus 3.5 MB per mailbox
Heavy	2 GB plus 5 MB per mailbox

Example:

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

- 2 GB + (3.5 MB x 4,000 mailboxes) = **16 GB RAM** (rounded to the nearest GB).
- 16 GB per server x 4 servers = **64 GB RAM**

4.4.3. Planning Mailbox Server Storage Requirements

Each Exchange server role has unique storage requirements with respect to throughput (IO) and capacity. Microsoft has detailed guidelines available on TechNet in the [Planning Storage Configurations](#) section of the Exchange Server 2007 documentation. Microsoft also provides the [Exchange Server 2007 Mailbox Server Role Storage Requirements Calculator](#) 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 2007.

Mailbox Server Role Storage Requirements Calculator v14.7 Inputs

Input	Value
Exchange Server 2007 version	SP1+
Number of Exchange mailbox servers	4
High availability model	None ²
Number of SCR targets/source server	0
Use recommended maximum database sizes?	Yes
Total number of tier-1 user mailboxes	16,000
Send/receive capability/mailbox/day	10 sent/40 received
Average message size (KB)	50
Tier-1 user mailbox size limit (MB)	500
Backup methodology	Software VSS backup/restore
Backup frequency	Weekly full/daily incremental

Using the Exchange Server 2007 Mailbox Server Role Storage Requirements Calculator version 14.7 and the above variables, our Mailbox server role storage configuration is summarized as follows:

Mailbox Server Role Storage Requirements Calculator v14.7 outputs

Output	Value
Database/log configuration (per server)	21 databases per server 120 GB database size + overhead 11 GB log size + overhead
Storage group LUN design (per server)	7 databases per LUN SG1-SG7: 1101 GB DB LUN/96 GB Log LUN SG8-SG14: 1101 GB DB LUN/96 GB Log LUN SG15-SG21: 1101 GB DB LUN/96 GB Log LUN
RAID configuration (per server)	Databases per server 20 x 400 GB 10,000 RPM FC/SCSI/SAS 3.5" Logs per server 4 x 300 GB 15,000 RPM FC/SCSI/SAS 3.5"
RAID configuration (total)	Databases (4 servers) 80 x 400 GB 10,000 RPM FC/SCSI/SAS 3.5" Logs (4 servers) 16 x 300 GB 15,000 RPM FC/SCSI/SAS 3.5"

²The high availability model should be chosen to meet your organization's specific business and service-level requirements. VMware supports all of the current application-aware high availability models for Microsoft Exchange Server 2007, including CCR clustering, SCC clustering, and LCR. VMware provides VMware HA and VMware VMotion as additional levels of protection, regardless of the Exchange mailbox server model.

4.4.4. Building Block Summary (4,000-User Average Profile)

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

4,000-User Average Profile Building Requirements

Exchange Role	Physical Resources (per server)
Mailbox server	CPU: 4 cores Memory: 16 GB OS and Application File Storage: 16 GB (OS & application files) Database Storage: 20 x 400 GB 10,000 RPM FC/SCSI/SAS 3.5" Log Storage: 4 x 300 GB 15,000 RPM FC/SCSI/SAS 3.5" (logs) Network: 1 Gbps

4.4.5. Guest Virtual Machine Configuration

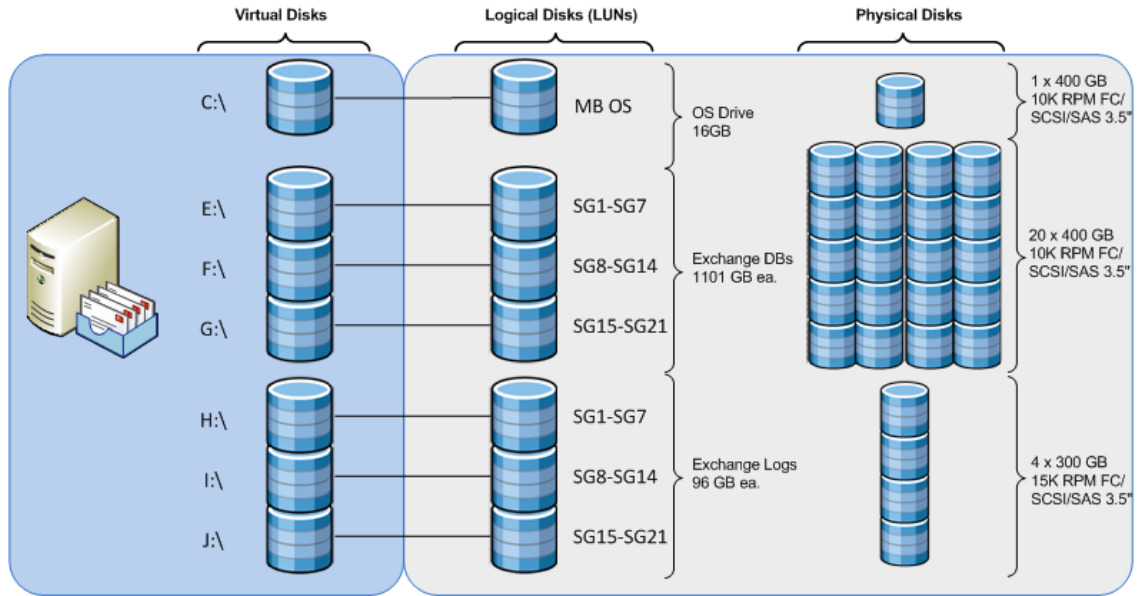
The resource requirements for Exchange are translated below into the virtual machine configuration for our 4,000-user mailbox server building block.

Exchange Virtual Machine Configuration

Exchange Role	Virtual Hardware (per VM)
Mailbox Server	CPU: 4 vCPU Memory: 16 GB Storage: SCSI Controller 0 HDD 1: 16 GB (OS & application files) HDD 2: 1101 GB (SG1-SG7 databases) HDD 3: 1101 GB (SG8-SG14 databases) HDD 4: 1101 GB (SG15-SG21 databases) HDD 5: 96 GB (SG1-SG7 logs) HDD 6: 96 GB (SG8-SG14 logs) HDD 7: 96 GB (SG15-SG21 logs) Network: NIC 1

4.4.6. Guest Virtual Machine Storage Interaction

The figure below demonstrates visually how the building block virtual machine will interact with the shared storage.



Exchange Mailbox VM
(4,000 user/average profile)

Shared Storage

4.5. Sizing Peripheral Server Roles

4.5.1. Server Role Ratios

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. The following table summarizes the guidelines Microsoft recommends:

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	4:1

4.5.2. CPU and Memory Requirements

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 sever role. The following table summarizes the guidelines that Microsoft recommends:

Memory configurations for Exchange Server 2007 servers based on installed server roles

Server role	Minimum per server	Recommended	Maximum per server
Hub Transport	2 GB	1 GB per core (2 GB minimum)	16 GB
Client Access	2 GB	2 GB per core (2 GB minimum)	16 GB

Example:

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

Hub Transport Calculations

16 mailbox processor cores/7 hub core ratio with AV = **4 hub processor cores** (rounded for even distribution)

4 processor cores x 1 GB per core = **4 GB RAM**

Based on your business requirements, you may decide to deploy one quad-core/4 GB RAM instance of the Hub Transport server or two, dual-core/2 GB RAM instances of the Hub Transport server.

Client Access Server Calculations

16 mailbox processor cores/4 Client Access core ratio = **4 client access processor cores**

4 processor cores x 2 GB per core = **8 GB RAM**

Based on your business requirements, you may decide to deploy one quad-core/8 GB RAM instance of the Client Access server or two, dual-core/4 GB RAM instances of the Client Access server.

4.6. Planning Guest Virtual Machine Configurations

With any application, you should plan its resource requirements first, and then the underlying infrastructure requirements to meet the applications needs.

4.6.1. Resource Requirements by Server Role

Example:

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

Example Exchange Server Role Resource Requirements

Exchange Role	Physical Resources (per server)
Mailbox Server (4 servers)	CPU 4 cores Memory 16 GB Storage 16 GB (OS & Application files) 20 x 400 GB 10,000 RPM FC/SCSI/SAS 3.5" (Databases) 4 x 300 GB 15,000 RPM FC/SCSI/SAS 3.5" (logs) Network 1 Gbps
Client Access Server (2 servers)	CPU 2 cores Memory 4 GB Storage 16 GB (OS & application files) Network 1 Gbps
Hub Transport Server (2 servers)	CPU 2 cores Memory 2 GB Storage 16 GB (OS, application, & log files) 32 GB (DB, protocol/tracking logs, & temp files) Network 1 Gbps

4.6.2. 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.

Example Exchange Virtual Machine Configuration

Exchange Role	Virtual Hardware (per VM)
Mailbox Server (4 VMs)	CPU 4 vCPU Memory

	16 GB Storage SCSI Controller 0 HDD 1: 16 GB (OS & application files) HDD 2: 1101 GB (SG1-SG7 databases) HDD 3: 1101 GB (SG8-SG14 databases) HDD 4: 1101 GB (SG15-SG21 databases) HDD 5: 96 GB (SG1-SG7 logs) HDD 6: 96 GB (SG8-SG14 logs) HDD 7: 96 GB (SG15-SG21 logs) Network NIC 1
Client Access Server (2 VMs)	CPU 2 vCPU Memory 4 GB Storage SCSI Controller 0 HDD 1: 16 GB (OS & application files) Network NIC 1
Hub Transport Server (2 VMs)	CPU 2 vCPU Memory 2 GB Storage SCSI Controller 0 HDD 1: 16 GB (OS & application files) HDD 2: 32 GB (DB, protocol/tracking logs, & temp files) Network NIC 1

4.7. Planning the Host Hardware Configuration

4.7.1. Virtual Machine Distribution

Now that we understand the physical resource requirements and associated virtual hardware configuration, we can plan physical VMware ESX host hardware to meet those requirements. To build infrastructure availability into the architecture, we will distribute the 8 total VMs across 2 physical VMware ESX host servers.

Example Exchange Virtual Machine Distribution

ESX Host	VM(s)
ESX Host 1	Exchange Mailbox VM 1 (4 vCPU/16GB RAM) Exchange Mailbox VM 2 (4 vCPU/16 GB RAM) Exchange Client Access VM 1 (2 vCPU/4 GB RAM) Exchange Hub Transport VM 1 (2 vCPU/2 GB RAM)
ESX Host 2	Exchange Mailbox VM 3 (4 vCPU/16GB RAM) Exchange Mailbox VM 4 (4 vCPU/16 GB RAM) Exchange Client Access VM 2 (2 vCPU/4 GB RAM) Exchange Hub Transport VM 2 (2 vCPU/2 GB RAM)

4.7.2. Host Specifications

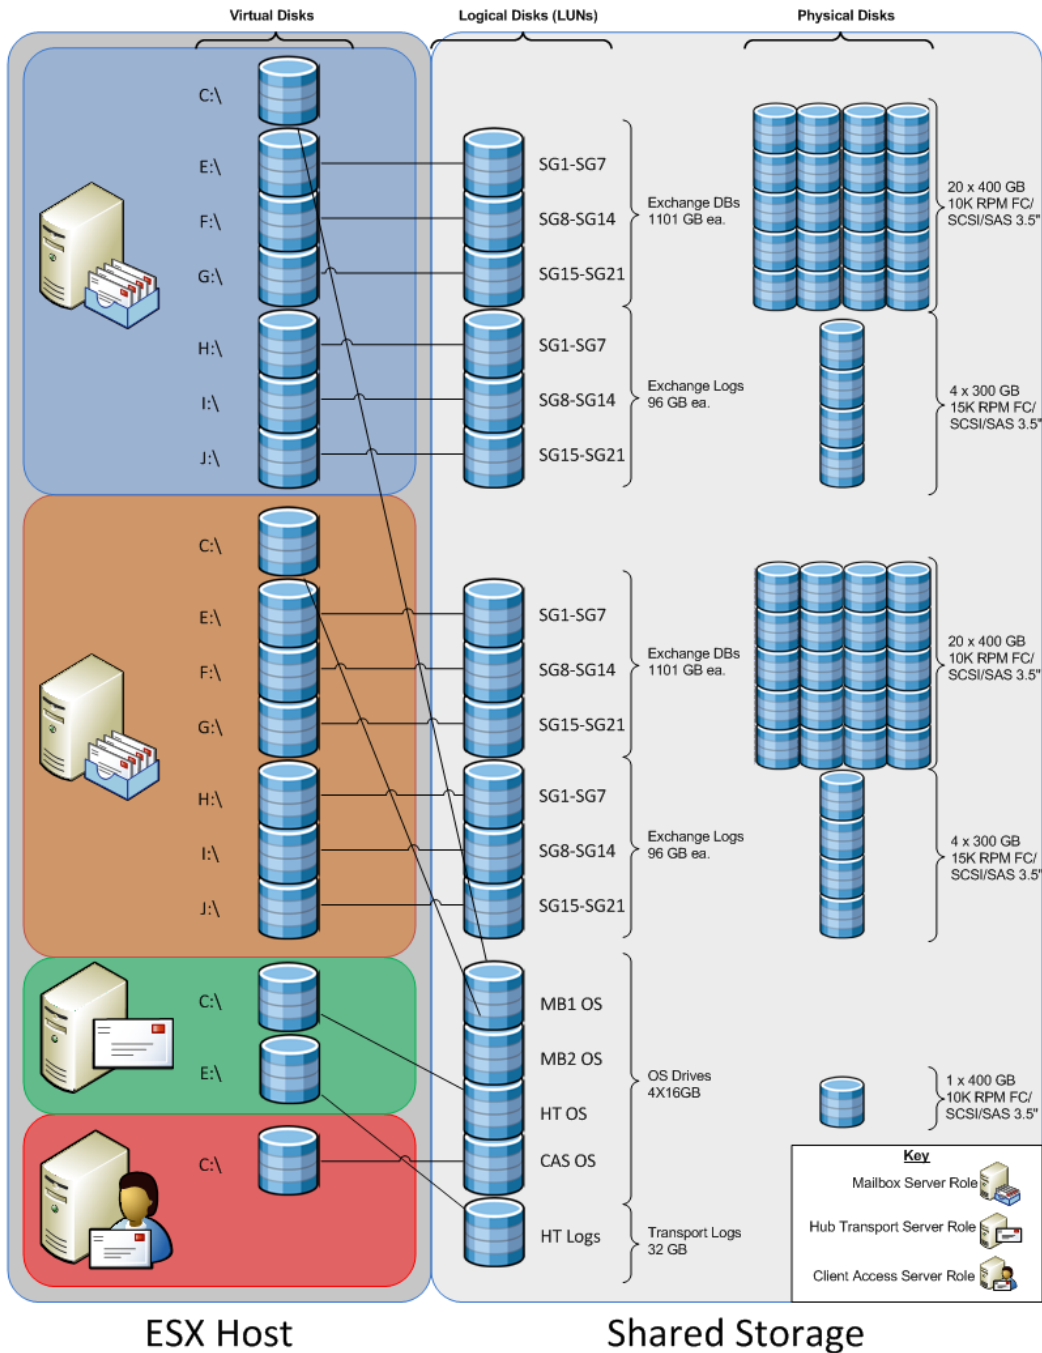
Each VMware ESX host server should provide enough physical hardware resources to accommodate the planned workload and provide some headroom in the event of a VMware HA failover or planned VMware VMotion migration of live virtual machines for host hardware maintenance. The following table summarizes the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

Example ESX Host Hardware Configuration Table

ESX Host	VM(s)
ESX Host 1	16 cores (4x4) 48 GB RAM 2 Fiber Channel HBAs 4 Gigabit network adapters
ESX Host 2	16 cores (4x4) 48 GB RAM 2 Fiber Channel HBAs 4 Gigabit network adapters

4.8. VMware ESX Host Architecture

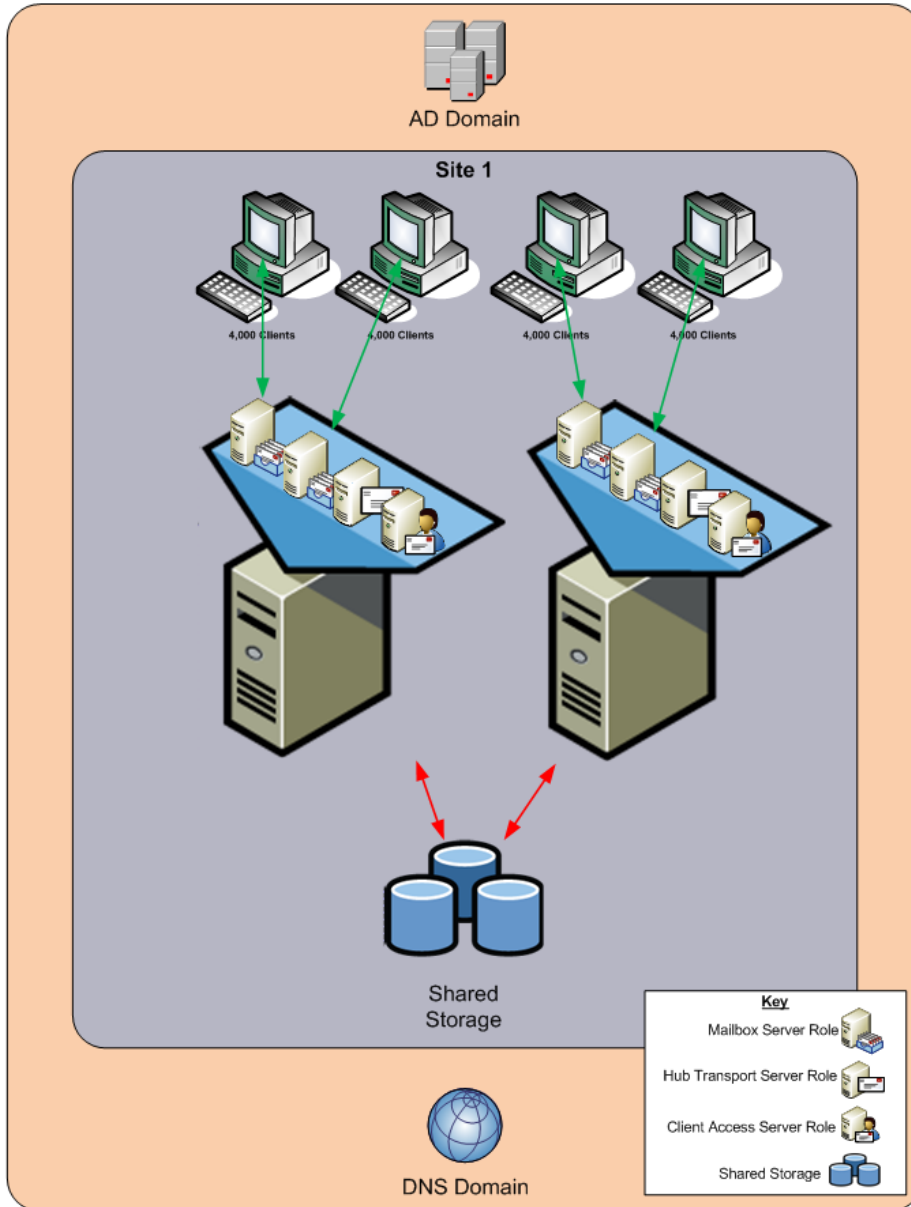
Based on the above design decisions, each VMware ESX host would run two Mailbox Server virtual machines, one Hub Transport virtual machine, and one Client Access Server virtual machine. The diagram below illustrates how the VMware ESX Host and each virtual machine interacts with the underlying storage:



VMware ESX Host Architecture for Exchange Server 2007 on VMware Infrastructure

4.9. Logical Architecture

The logical architecture is a basic single-site, single-domain design, illustrated in the figure below. Each Exchange Mailbox Server virtual machine is equipped to handle 4,000 mailboxes. The entire solution is designed to support 16,000 total users.



Logical Architecture for Exchange Server 2007 on VMware Infrastructure

5. Deployment and Operations

You can leverage VMware Infrastructure 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.

5.1. VMware Infrastructure Deployment Features

5.1.1. 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 pre-configured Exchange system. In native environments, this process can consume significant time, requiring you to procure hardware and install the operating system. Cloning ensures a controlled virtual machine configuration so deployment is less error-prone and less time-consuming.

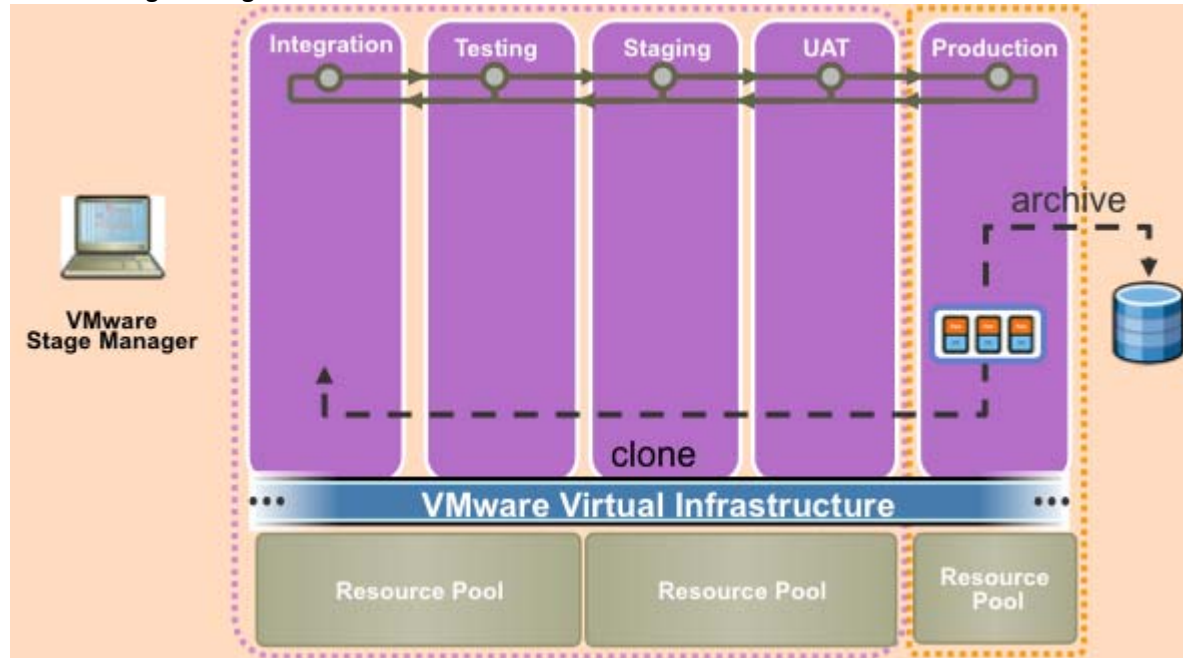
5.1.2. VMware Stage Manager

VMware Stage Manager can streamline the testing and deployment of configuration changes, patches, or upgrades to your Exchange infrastructure. When you need to make changes to the production Exchange systems, Stage Manager allows you to take a snapshot of the current Exchange service configuration into an archive.

Archiving is an explicit function in VMware Stage Manager. It is also very easy to restore service configurations from the archive, which is helpful if you need to restore historical snapshots for an audit when hardware configurations, applications, or database schemas have changed. You can then clone a production system, deploy it in a pre-production resource pool, apply your patches, and transition the patched system through the discrete stages of a release management cycle. If the patch is tested successfully, you can either promote the patched system into production, replacing the previous copy, or apply the patches directly to the current production system.

The service configuration clone is an exact replica of the Exchange production system, including all network settings, host names, and IP addresses. VMware Stage Manager deploys the clone in a “fenced network” to prevent network collisions. The “fenced networking” feature allows simultaneous deployment of multiple instances of the exact same Exchange service configuration, which allows multiple teams to work in parallel without interrupting or conflicting with one another.

VMware Stage Manager



5.2. VMware VMotion, VMware DRS, and VMware HA

VMware 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 VMware VMotion capability a step further by adding an intelligent scheduler. DRS allows you to set resource assignment policies that reflect business needs. VMware 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 VMware DRS perform best under the following conditions:

- The source and target VMware ESX Server hosts must be connected to the same gigabit network and the same shared storage.
- A dedicated gigabit network for VMware VMotion is recommended.
- The destination host must have enough resources.
- The virtual machine must not use physical devices like 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 (e.g. 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: VMware does not currently support VMware VMotion or VMware DRS for Microsoft Cluster nodes; however, a cold migration is possible once the guest OS is shut down properly.**

With VMware High Availability (HA), Exchange virtual machines on a failed VMware ESX Server host can be restarted on another VMware ESX Server 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 VMware ESX Server host hardware failure and does not monitor the status of the Exchange services—these must be monitored separately.
- Proper DNS hostname resolution is required for each VMware ESX Server host in a VMware HA cluster.
- VMware HA “heartbeat” is sent via the VMware ESX Server service console network so redundancy in this network is recommended.



VMware, Inc. 3401 Hillview Ave Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Copyright © 2008 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,961,806, 6,961,941, 6,880,022, 6,397,242, 6,496,847, 6,704,925, 6,496,847, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,944,699, 7,069,413, 7,082,598, 7,089,377, 7,111,086, 7,111,145, 7,117,481, 7,149,843, 7,155,558, 7,222,221, 7,260,815, 7,260,820, 7,268,683, 7,275,136, 7,277,998, 7,277,999, 7,278,030, 7,281,102, 7,290,253; patents pending.

