Microsoft Exchange 2013 on VMware Design and Sizing Guide

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

Contents

1. Introduction 5
2. Design Concepts 6
   2.1 Data Gathering 6
   2.2 Building the Functional Design 7
   2.3 Defining Compute Requirements 8
   2.4 Application of Compute Requirements to the Virtual Platform 16
   2.5 Establishing Virtual Machine Sizing and Placement 17
   2.6 Sample Physical Layout 18
3. Sizing Examples 19
   3.1 Single Role Server Design (12,000 Users) 19
   3.2 Single Role Server Design with DAG for 24,000 Users 23
   3.3 Multirole and Multisite DAG Server Design (50,000 Users) 27
4. Summary 34

List of Tables

Table 1. Megacycles per Mailbox 9
Table 2. Minimum Memory Requirements 12
Table 3. Per Mailbox Database Cache 13
Table 4. Determining Total Memory 14
Table 5. Exchange Server Role Resource Requirements 19
Table 6. Exchange Virtual Machine Configuration 20
Table 7. Exchange Virtual Machine Distribution 21
Table 8. ESXi Host Hardware Configuration Table 22
Table 9. Exchange Server Role Resource Requirements 23
Table 10. Exchange Virtual Machine Configuration 24
Table 11. Exchange Virtual Machine Distribution 26
Table 12. ESXi Host Hardware Configuration Table 26
Table 13. Exchange Server Role Resource Requirements 28
Table 14. Exchange Virtual Machine Configuration 29
Table 15. Exchange Virtual Machine Distribution: Datacenter A 31
Table 16. Exchange Virtual Machine Distribution: Datacenter B 32
Table 17. ESXi Host Hardware Configuration Table 32

List of Figures

Figure 1. Sample Physical Layout for 24,000 Mailboxes 18
Figure 2. Mailbox Virtual Machine Configuration 21
Figure 3. Initial Virtual Machine Placement 22
Figure 4. Mailbox Virtual Machine Configuration 25
Figure 5. Initial Virtual Machine Placement 27
Figure 6. Mailbox Virtual Machine Configuration 31
Figure 7. Initial Virtual Machine Placement 33

1. Introduction

Microsoft Exchange can be complex to deploy, and there are many design decisions to make to build a solid solution. Running Microsoft Exchange Server 2013 on VMware® vSphere® can positively impact design, deployment, availability, and operations, but what does such a solution look like?

This document explores sample architecture designs that illustrate Exchange 2013 environments running on vSphere. The focus of this architecture is to provide a high-level overview of the solution components, with diagrams to help illustrate key concepts. For detailed best practices, see the *Microsoft Exchange 2013 on VMware Best Practices Guide*.

This design and sizing guide covers:

- Design concepts.
  - Data gathering.
  - Building the functional design.
  - Defining compute requirements.
  - Applying the compute requirements to the virtual platform.
  - Establishing virtual machine sizing and placement.
- Sizing examples.
  - Single role server design – 12,000 users.
  - Single role server design with DAG – 24,000 users.
  - Multisite design with multirole servers and DAG – 50,000 users.
- Design and deployment considerations.

The examples show how these components contribute to the overall design and provide only a guideline. Customers should work with their infrastructure vendors to develop a detailed sizing and architecture plan designed for their requirements. After describing some design concepts, this document looks at sizing examples of Exchange 2013 on vSphere using various design options and explores options using standalone mailbox servers and database availability group (DAG) servers using scale-out and multirole deployment methods.

This document provides examples to help understand components and concepts. Official sizing for Exchange environments varies based on business and technical requirements, as well as server and storage hardware platforms. VMware recommends that you engage your server and storage vendors to help plan your design, or use one of the detailed, hardware-specific reference architectures found on the VMware Web site and in the *Microsoft Exchange 2013 on VMware Partner Resource Catalog*.

2. Design Concepts

One of the most common questions about the virtualization of Exchange Server is regarding design and sizing. There is often the misconception that designing Exchange for running on vSphere requires special tools, a different approach, or vast knowledge of virtualization. In fact, many of the successful Exchange virtualization projects that VMware has delivered have been based on existing Exchange designs that were originally created based on a physical server deployment. The logical Exchange design is not impacted significantly by virtualization. Sizing, virtual machine placement, and how best to use features of vSphere ultimately drive what the Exchange topology looks like from a server count and distribution perspective.

Designing a new environment to support a virtualized Exchange environment follows the same basic process for a non-virtualized deployment, with a few additional steps. At a high level, the process includes:

- Data gathering.
- Building the functional design.
- Defining compute requirements.
- Application of the compute requirements to the virtual platform.
- Establishing virtual machine sizing and placement.

The following sections look at what is involved during each of these phases.

2.1 Data Gathering

Much of the input for the Exchange design comes from the prerequisite data collected in this phase. This includes the following topics:

- Understanding business and technical requirements.
- Evaluating the current workload.
- Evaluating the health of the surrounding infrastructure.
- Understanding support and licensing considerations.

The data acquired from these prerequisites drives the functional design and helps to achieve the virtualization design.

2.1.1 Understanding Business and Technical Requirements

A clear understanding of the business requirements for Exchange helps to drive much of the design. During this stage, questions about uptime requirements, growth expectations, feature support, security, and regulatory compliance requirements are answered. Many of these requirements are then mapped to specific features that can be provided either by Exchange itself or in combination with vSphere. For example, in the case of security, an organization might require that the email system be isolated from other applications within the datacenter. VMware vCloud® Networking and Security™ can help to achieve application isolation.

It should be noted that in some organizations virtualization takes priority in design consideration. What this means is that the application design must conform to what the virtualized infrastructure can provide. For example, VMware vSphere High Availability (HA) should be used as the primary method of high availability instead of an application-specific clustering solution. This falls into the business and technical requirements discussion.

2.1.2 Evaluate the Current Workload

In most environments, with the exception of very new organizations, an established Exchange or other email environment is used to evaluate the workload characteristics of users. Microsoft bases guidance for Exchange server sizing requirements on the activity of users. This includes how many messages are sent and received per day, whether or not additional client types are used (such as mobile devices and archiving systems), and average message size. Much of this data can be collected, using native tools such as Microsoft Perfmon, by gathering and parsing mail log files or by using third-party tools. Regardless of the method used, an understanding of the type of load the user base will put on the proposed environment is an absolute requirement to sizing Exchange properly.

There are scenarios when user characteristics are not known or cannot be evaluated. This is the case in a new environment, such as a new company. For Exchange there is good data as to what performance characteristics will be like depending on the number of messages sent and received per day. A good starting point for most environments, even those with established workload characteristics, has been the *150 messages sent and received per day* user profile. Although not very scientific, it provides a safe starting point for most organizations, especially new ones. Because vSphere is a very flexible platform, you can scale up or out, or even down, as needed.

2.1.3 Evaluate the Health of the Surrounding Infrastructure

Exchange is highly dependent on services provided by Active Directory, DNS, the network infrastructure, and the storage area network (SAN), assuming that storage is based on SAN technology. Although everyday user activity, such as authentication and name resolution, might appear to function as intended, the introduction of an application such as Exchange can make deficiencies in the infrastructure much more apparent. A thorough health check of the infrastructure should be performed before any Exchange software is installed because even the installation of Exchange is dependent on Active Directory being completely functional. A faulty Active Directory can cause the installation of Exchange to fail and lead to support calls followed by hours of manual cleanup.

2.1.4 Understand Support and Licensing Considerations

Although support and licensing is not an area of much concern for Exchange 2013, it is important to be familiar with this topic. Exchange 2013 is fully supported on vSphere as a result of the Windows Server Virtualization Validation Program. However, there are certain caveats to support, such as the level of CPU overcommitment supported for production environments and the use of network-attached storage (NAS).

Licensing is typically a straightforward conversation with Exchange because of the continued use of a client/server licensing model. Other applications are not as simple, such as Microsoft SQL Server. It is important to understand the licensing implications that might affect design decisions, such as scaling out versus scaling up. Refer to the *Microsoft Exchange 2013 on VMware Support and Licensing Guide* for more information.

2.2 Building the Functional Design

Knowledge of Exchange 2013 architecture is required during this phase of the design. At the most basic level you can build an Exchange design, and in most cases, translate that directly to virtual machines and have a functional Exchange environment. However, knowledge of vSphere, its configuration options, such as the number of vCPUs and disk targets supported per virtual machine, and design best practices allow an Exchange architect to make the best decisions for a virtualized Exchange environment.

A successful Exchange virtualization project should begin with all areas of the infrastructure represented in the conversation. This includes Exchange, vSphere, storage, networking, and any other areas for consideration, such as facilities. During the design discussions, each functional area is discussed, and input from the various technical representatives is collected for further consideration.

The design requirements that result in the functional design comprise both hard and soft values. Items such as the number of mailboxes, user profile, the number of datacenters, and the tiers of mailboxes to support are hard values. These values have no additional options for consideration. Design requirements such as uptime, database size, and hardware specifications must be discussed further to determine the best option to meet the needs of the organization. Before some of these decisions are made there might be further testing required by the organization to validate the solution. In most cases the expertise of the architecture team should be able to speak to each option and its capabilities.

When completed, the functional design should include at least the following:

- High availability method – vSphere High Availability or Exchange DAG.
- Site resiliency method – None, VMware vCenter™ Site Recovery Manager™, or Exchange DAG.
- Dedicated or multirole servers.
- Database sizing.
- Data protection – Exchange Native Data Protection, Exchange-aware backup, VMware vSphere Data Protection Advanced™.
- Estimated growth over how many years.
- Hardware options – Preferred server vendor and deployment options, such as blade versus rack mount.
- Mailbox tiers – Number of mailboxes, mailbox size or sizes, average message size, archive limit.
- Client connectivity – VMware vCloud Network and Security Edge™, hardware or software load balancer, Windows Network Load Balancing (NLB), or DNS round robin.

2.3 Defining Compute Requirements

With the functional design complete, the basics for understanding the physical compute requirements are established. To begin the process of defining the compute requirements, there must be an understanding of what is involved in this process.

The official sizing guidance for Exchange 2013 was not available from Microsoft at the time of this writing. However, with the consolidation of server roles in Exchange 2013, the Mailbox server role has become much like the Exchange 2010 multirole client access, hub transport, and mailbox server. This section reviews the process for defining compute requirements. During the discussion of defining compute requirements, examples are provided to help illustrate consistently the main points discussed.

Note Although the values used in these examples are specific to Exchange 2010, the methodology remains the same. As Microsoft provides updated guidance for Exchange 2013, replace the following values with updated values, if necessary.

Example

The following examples look at the basic sizing of an Exchange 2013 environment. These values are used to determine the compute requirements, sizing, and placement of virtual machines for this environment.

- Total mailboxes – 24,000
- Average mailbox quota – 2048MB
- Average daily send/receive – 150 messages
- Average message size of 75KB
- High availability– database availability group, vSphere HA
- Database copies – 2
- Sites – one site
- Processor architecture – Eight-core processor with a SPECint2006 rating of 41 per core

2.3.1 Processor Core Requirements

CPU requirements for Exchange mailbox servers are represented in megacycles. A megacycle is a unit of measurement used to represent the capacity of a processor core. The performance delivered by a processor core is defined by the clock speed of the processor core. For example, a 3.33GHz processor core provides 3,333 megacycles. This is the baseline used by Microsoft to provide guidance for the megacycle requirement of a mailbox profile. The following table provides the megacycle estimates for various mailbox profiles for Exchange 2010. Until further guidance for Exchange 2013 is provided by Microsoft, these numbers should continue to be used as a starting point.

Table 1. Megacycles per Mailbox

Messages Sent or Received per Mailbox per Day	Megacycles for Active Mailbox or Standalone 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

With the advancement of processor technology, simply using the megacycles provided by a processor core is no longer adequate. Many newer processor cores operate at a lower clock speed than the baseline used by Microsoft but provide higher throughput. As a result a megacycle adjustment is required to determine the actual capabilities of a processor core. To make this adjustment, processor throughput ratings from the Standard Performance Evaluation Corporation (SPEC) are used to determine the difference between the baseline per core value and the per core value of a newer processor. SPECint2006 rate results for processors are found on the SPEC website using the search feature (http://www.spec.org/cgi-bin/osgresults?conf=rint2006).

As an example, the baseline processor used by Microsoft, the Intel Xeon X5470 (3.33GHz), has a SPECint2006 rating of 18.75 per core. The Intel Xeon E5-2670 (2.60GHz) eight-core processor has a SPECint2006 rating of 41 per core. This is roughly a 218% improvement. To calculate the adjusted performance per core of the new processor, use the following formula.

((new per core value) * (baseline Hertz per core)) / (baseline per core value) = adjusted megacycles per core

Using this example:

((41) * (3333) / 18.75 = 7288 adjusted megacycles per core

This value is used to determine how many mailboxes can be supported on a given processor core. Exchange workloads should maintain a one-to-one physical processor core to virtual CPU ratio. This allows for a true representation of adjusted megacycle capabilities when designing and deploying Exchange on vSphere.

2.3.1.1. Calculating the Megacycle Requirement for Standalone Mailbox Servers

The process for calculating the megacycle requirement for a mailbox server can take two forms depending on whether a DAG is used. Standalone mailbox servers, servers not in a DAG, must provide resources only for the mailboxes that they are going to support during normal runtime. In other words, if an average mailbox consumes 3 megacycles, and the mailbox server must support 2000 mailboxes, the mailbox server must be able to deliver 6000 megacycles of processor capacity. To provide for the occasional spike in utilization, Microsoft typically recommends establishing a maximum utilization threshold. In Exchange 2010 the threshold for a standalone mailbox server with all roles installed was 35%. This can be used as a baseline for Exchange 2013.

The following summarizes the process to determine the megacycle requirement for a standalone mailbox server supporting 2000 users at 3 megacycles per user.

1. Determine the total mailbox megacycle requirements – 2000 mailboxes * 3 megacycles/user = 6000 megacycles.
2. Adjust megacycles for 35% peak utilization – 6000 megacycles / .35 = 17142 total megacycles required.

Using processor cores from the previous example that support 7,288 adjusted megacycles per core, a mailbox server with two cores utilizes approximately 40% of its CPU capacity. This is more than the recommended threshold of 35%, but given the option of overprovisioning the virtual machine by adding an additional core or two, this is an acceptable configuration with plenty of capacity for spikes.

2.3.1.2. Calculating Megacycle Requirement for DAG Member Servers

DAG member servers require a more in-depth megacycle calculation process because of variables such as the maximum number of active mailboxes per server, the number of passive mailboxes per server, and the number of database copy instances. In Exchange 2010 the threshold for a DAG member server with all roles installed was 40%. This can be used as a baseline for Exchange 2013.

The following summarizes the process to determine the megacycle requirements for a DAG member server in a four-node DAG supporting 16,000 users at 3 megacycles per user. Two copies per database are used in this example, allowing for one DAG member server failure. This example uses the same processor cores as in the preceding section.

1. Determine the maximum active mailboxes per DAG member server after a single server failure – 16000 total mailboxes / (4 DAG members - 1) = 5334 maximum users per DAG member server after a single server failure.
2. Determine the active mailbox megacycle requirements – 5334 maximum users per DAG member server * 3 megacycles per active user = 16002 megacycles.
3. Add 10% per additional database copy—in this example there is one additional copy – 16002 megacycles * 1.1 = 17602 megacycles to support active mailboxes per DAG member server.
4. Determine the number of passive mailboxes per DAG member server after a single server failure – (4000 active mailboxes during normal runtime * 2 database copies) - (5334 maximum active mailboxes) = 2666 passive mailboxes.
5. Determine the passive mailbox megacycle requirements – 2666 passive mailboxes * 0.45 megacycles/passive mailbox = 1200 megacycles.

   Note The passive mailbox megacycle requirement is about 15% of the active mailbox requirement. See Table 1 for a complete listing.

6. Add active and passive mailbox megacycle requirements to determine total megacycle requirements per DAG member server – 17602 active mailbox megacycles + 1200 passive mailbox megacycles = 18802 megacycles.
7. Adjust megacycles for 40% peak utilization – 18802 total megacycles / 0.4 = 47005 megacycles required per DAG member server.
8. Determine the number of processor cores required – 47005 total megacycles / 7288 adjusted megacycles per core = 6.4 processor cores.

Using processor cores from the previous example, which support 7,288 adjusted megacycles per core, a mailbox server with six cores utilizes approximately 43% of its CPU capacity. This is over the recommended threshold of 40%, but given the option of overprovisioning to seven cores, this is an acceptable solution.

2.3.1.3. Processor Sizing Considerations

The preceding example shows an established starting point for the number of mailbox server virtual machines to begin the sizing exercise. The megacycles calculation generates the megacycle requirement based on the number of virtual machines. This provides an idea of exactly how much CPU performance is required. Different virtual machine counts yield different results and should always be considered to find the best balance between the supported number of mailboxes per instance, virtual machine size, and number of virtual machines.

When sizing virtual machines from a CPU perspective, consider the following:

- vSphere virtual machines support up to 64 virtual CPUs.
- Microsoft typically recommends a minimum and maximum supported number of CPU cores. For Exchange 2010 the maximum was 12 cores for single-role mailbox servers and 24 cores for multirole servers.
- Microsoft recommends using a hypervisor overhead of 10% when calculating CPU requirements. VMware has seen hypervisor overhead as low as 2% for Exchange workloads on the latest vSphere hypervisor. In most cases, using the latest vSphere version and processors help to mitigate any hypervisor overhead.

Example

This example shows support for 24,000 users, protected by DAG and with a mailbox profile of 150 messages sent/received per day. Because the DAG supports two mailbox database copies, you must begin with a multiple of two for the number of DAG member servers. Two servers is the minimum, but you can scale out from there to accommodate any deployment scenario. This example assumes 4 DAG nodes. Adjustments can be made later, if desired.

To calculate the megacycle requirements, perform the following:

- 24,000 / 4 DAG members = 6,000 mailboxes per DAG member server during normal operations.
- 24,000 mailboxes / (4 DAG members - 1) = 8,000 maximum mailboxes per DAG member server.
- 8,000 maximum mailboxes * 3 megacycles per active mailbox = 24,000 megacycles required.
- 24,000 * 1.1 to account for the additional database copy = 26,400 megacycles.
- (6,000 mailboxes during normal operations * 2 database copies) - (8,000 maximum mailboxes per DAG member server) = 4,000 passive mailboxes.
- 4,000 passive mailboxes * 0.45 megacycles = 1800 passive mailbox megacycles.
- 26,400 active mailbox megacycles + 1,800 passive mailbox megacycles = 28,200 total megacycles.
- 28,200 total megacycles / 0.40 maximum CPU utilization during failover = 70,500 megacycles required per DAG member server.
- Each proposed processor core has a SPECint2006 rating of 41, and provides 7,288 adjusted megacycles.
- 70,500 megacycles / 7,288 megacycles per core = 10 cores per DAG member server.

At 10 cores, or virtual CPUs, each DAG member server is approximately 40% utilized after a single DAG member server failure. The number of DAG member servers can be scaled out even further if smaller virtual machines are desired.

2.3.2 Memory Requirements

Proper sizing of memory resources is much less complicated than processor sizing. The amount of memory assigned to an Exchange 2013 mailbox server depends on the maximum active user count to be supported on the mailbox server and the profile of those mailboxes. This provides the database cache for user data. Additional memory must be provided to support the operating system and other applications.

Exchange 2013 does have minimum memory support requirements. The following table shows these minimums.

Table 2. Minimum Memory Requirements

Exchange 2010 Server Role	Minimum Supported
Client Access	4GB
Mailbox	8GB
Client Access and Mailbox combined	8GB

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, to support 4,000 users sending/receiving 150 messages per day requires 36GB of database cache using the Exchange 2010 recommendation of 9MB of database cache per mailbox (4000 * 9MB = 36GB).

The following table shows the recommended per mailbox database cache used to size Exchange 2010 mailbox servers. Initial sizing of Exchange 2013 environments can continue to use these numbers until official guidance from Microsoft is released.

Table 3. Per Mailbox Database Cache

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

The next step is to determine the amount of required physical memory by determining which server configuration provides enough database cache, as well as additional memory, for the operating system and applications. Microsoft has provided examples of common memory configurations and how much database cache would be provided with that configuration. Current guidance provided is specific to Exchange 2010, however based on the architecture changes in Exchange 2013, sizing guidance for Exchange 2010 multirole mailbox servers can be used as a starting point for Exchange 2013 mailbox servers.

The preceding example shows that 4,000 users sending and receiving 150 messages per day requires 36GB of database cache. Based on the following table, a mailbox server with 64GB of physical RAM provides 44GB of database cache. Therefore, 64GB of physical RAM is the ideal memory configuration, based on this mailbox count and user profile.

Table 4. Determining Total Memory

Server Physical Memory	Database Cache Provided
8GB	2GB
16GB	8GB
24GB	14GB
32GB	20GB
48GB	32GB
64GB	44GB
96GB	68GB
128GB	92GB

Example

In this example each DAG member server supports a maximum of 8,000 mailboxes after a single server failure. To calculate the minimum recommended database cache per DAG member server, perform the following:

- 8,000 maximum active mailboxes per DAG member server * 9MB of database cache per user = 72GB
- 72GB of database cache is required to support 8,000 active mailboxes, and additional memory is required to support the operating system and applications.
- According to Table 4, to provide 72GB of database cache, each DAG member server should be allocated 128GB of memory.

Each DAG member server virtual machine is created with 128GB of memory allocated.

2.3.3 Network Requirements

Exchange virtual machines configured with the Client Access or Mailbox server role, or both, and not participating in a DAG, typically require no more than a single virtual network adapter. When deployed in a DAG, a virtual machine can be configured with a single network adapter, but the recommended configuration for DAG nodes is to provide a network adapter for client communication and a separate adapter for DAG replication. Within a virtual machine this means configuring two virtual network adapters and connecting those adapters to port groups or virtual switches dedicated to each type of traffic.

At the VMware ESXi™ host level, a minimum of two physical network adapters should be teamed for redundancy and configured based on VMware best practices. Separate VLANs are recommended to separate vSphere management traffic, as well as client and DAG replication traffic, for Exchange virtual machines. Refer to the *Microsoft Exchange 2013 on VMware Best Practices Guide* for more information.

2.3.4 Storage Requirements

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 enables more accurate planning when migrating from email systems other than Microsoft Exchange. The user profile has a direct impact on overall I/O requirements, and knowing these requirements can help you and your storage vendors to design an optimal storage solution. In addition to the average mail sent and received, mobile devices, archiving solutions, and antivirus programs should be taken into consideration as contributors to overall I/O.

Exchange 2013 continues with the reduction in I/O, making more storage options available. Some of the new features in Exchange 2013 include the support for multiple databases per disk and automatic reseed of databases with automated disk recovery. These new features are geared towards environments deploying on the larger, more failure-prone disk drives with no RAID-level storage redundancy. The typical deployment scenario when using *Just a Bunch of Disks* (JBOD) includes managing three or more copies per database because of the likelihood of failure and requirement to reseed in the case of a single disk failure. This in turn increases management overhead.

Using local storage for virtual machines is supported for use with vSphere, however many customers continue to deploy on shared storage. When used with shared-storage architecture, vSphere provides access to all advanced features, such as vSphere HA, VMware vSphere Distributed Resource Scheduler™ (DRS), and VMware vSphere vMotion®. Using data protection mechanisms provided by most storage arrays allows for minimal database copy maintenance—most environments deployed on shared-storage deploy a maximum of two database copies for high availability. If a disk failure does occur, the data stored on the volume is not lost, and no reseed is required, assuming storage vendor best practices are followed.

Sizing storage for a virtualized Exchange environment is the same as sizing for a physical environment regarding I/O requirements. There are a few vSphere-specific items to consider when designing storage for the Mailbox server role, as follows:

- ESXi hosts can have up to 255 individual storage LUNs mapped to them. This should be considered a vSphere cluster maximum because the best practice is that all hosts in a vSphere cluster are mapped to the same storage. If more than 255 storage LUNs must be presented to all of your Exchange virtual machines, consider creating more, smaller vSphere clusters, consolidating virtual disks on larger VMware vSphere VMFS volumes, using larger raw device mappings, or using in-guest attached iSCSI.
- When using VMFS datastores for Exchange data, be aware that the default configuration of an ESXi host limits the open virtual disk capacity to 8TB. For more information, see *ESXi/ESX host reports VMFS heap warnings when hosting virtual machines that collectively use 4 TB or 20 TB of virtual disk storage* (http://kb.vmware.com/kb/1004424). This limit does not apply to raw device mappings or storage mapped using in-guest attached iSCSI.
- Up to 60 storage targets can be configured per virtual machine.
- Virtual machine disk format (VMDK) disks can be created up to 2TB. For larger volumes, physical-mode raw device mappings can be used up to 64TB and in-guest attached iSCSI can be used up to the guest operating system limit.

Microsoft has stated that storage sizing for Exchange 2013 is very similar to that of Exchange 2010. To assist in planning the storage design of the Mailbox server role, customers should continue to use the *Exchange 2010 Mailbox Server Role Requirements Calculator* (http://blogs.technet.com/b/exchange/archive/2009/11/09/3408737.aspx) until an Exchange 2013 equivalent is released. VMware recommends that Exchange architects follow Microsoft best practices along with the storage vendor’s best practices to achieve an optimal storage configuration for Exchange Server 2013.

2.3.5 Exchange Mailbox Server Role Requirements Calculator

Although most Exchange architects understand the process for identifying Exchange compute requirements, some may not perform the manual steps outlined in the preceding sections. The Exchange Mailbox Server Role Requirements Calculator has taken all of the processes discussed here and incorporated best practices, storage guidance, and more into a single Excel workbook. Using the calculator is the recommended method for sizing Exchange, even for virtual deployments. At the time of this writing the calculator is only available for Exchange 2010, however Microsoft has stated that sizing for the Exchange 2013 Mailbox server role will be similar to sizing for multirole Exchange 2010 servers—Mailbox, Client Access and hub transport server roles in a single Exchange instance.

For more information on the Exchange Mailbox Server Role Requirements Calculator, go to the Exchange team Web site at http://blogs.technet.com/b/exchange/archive/2009/11/09/3408737.aspx.

Example

There are many options for Exchange storage architecture. The Exchange Mailbox Server Role Requirements Calculator takes input, much of which is presented in these examples, and provides a database layout recommendation and storage allocation scheme. Using this example, the calculator presents the following option:

- 18 databases per server – active and passive.
  - 1.8TB maximum database size.
  - 88GB maximum log size.
- 18 volumes are created to house both database and logs on the same volume.
  - 2.5TB database + log volume size.
- 2 volumes are created for operating system and application storage.
- To accommodate this amount of storage (approximately 46TB per mailbox virtual machine), raw device mappings are used. This also allows the use of volumes greater than 2TB.

Note Exchange 2013 supports a maximum of 50 mounted databases. Consider this when using the Exchange 2010 calculator to obtain early adopter sizing guidance.

2.4 Application of Compute Requirements to the Virtual Platform

With compute requirements for all Exchange virtual machines established, the data is then converted into a set of physical requirements. For Exchange and other business critical applications this is a very trivial exercise. VMware recommends a one-to-one physical to virtual ratio when allocating compute resources. For Exchange, this is a very important point to communicate.

The expectation is often that virtualization is a way of consuming more resources than are available in a physical server. Although this might appear to be the case because of the ability to overcommit virtual CPUs and memory, what is actually happening in the hypervisor is an advanced sharing algorithm that allows each virtual machine to believe that it has dedicated resources. As the ratio of virtual CPUs to physical CPU cores grows, the hypervisor must schedule more requests across the finite physical resources. This can result in higher wait times for tasks that are ready to be scheduled. For the majority of workloads that are not very intensive, this is not a problem.

Exchange is a resource-intensive workload, and when sized correctly, consumes the CPU and memory provided very efficiently (to a certain degree). How Exchange uses its compute resources leaves little room for any added latency due to overcommitment of physical compute resources. This does not mean that resources cannot be overcommitted for an ESXi host running Exchange workloads, but there should be an established baseline before any overcommitment is introduced into the environment so that if any additional latency is observed, there is a baseline for comparison.

The compute requirements that have been established with the functional design are aggregated and mapped to physical CPU cores and memory. This provides the total amount of physical compute resources required across the vSphere cluster. In some cases, especially in smaller environments, the compute requirements might be small enough to come from one or two physical servers. The next section looks at why it is important to understand sizing and placement in addition to the compute requirements. Although the minimum compute requirements might fit into a small number of physical servers, the high availability design might lead to a scaled-out approach.

Example

The following information is determined by using the data generated throughout the examples:

- Deploying four Exchange mailbox virtual machines requires that each virtual machine is configured with 10 vCPUs and 128GB of memory.
- Each mailbox server virtual machine has 20 storage targets.
- The vSphere cluster must provide the following:
    - At least 40 physical CPU cores.
    - 512GB of memory.

2.5 Establishing Virtual Machine Sizing and Placement

Physical server capabilities have far surpassed how applications can effectively use those resources. This is often the reason why an organization looks to virtualization. Exchange is no exception. Although Exchange does a very good job of using resources efficiently, with many of the current processor architectures, efficient use requires either placing a very large number of users on a single instance of Exchange or modifying the server build to provide fewer resources. The first option means that any service interruption affects a larger number of users, and the second option means more datacenter resources are consumed for less return on the investment.

In a virtualized Exchange environment these problems are solved by creating virtual machines sized to meet the various requirements, taking manageability and resource utilization into consideration. For larger deployments, creating a smaller number of larger virtual machines enables good consolidation of mailboxes and room for other peripheral virtual machines. Small to mid-sized environments might prefer to scale out the design with smaller virtual machines, allowing them to support multiple workloads alongside Exchange.

The approach taken at this phase of the design depends on factors such as physical server sizing, high availability requirements, and mailbox count. Larger physical servers can accommodate larger mailbox server virtual machines without overcommitting physical resources, but with smaller physical servers, the design approach might use smaller virtual machine sizes. If a DAG is being considered to provide high availability, the best practice recommendation is to host one DAG member virtual machine per physical host. This means that the vSphere cluster must contain, at a minimum, the same number of physical hosts as proposed DAG nodes. If the design calls for multiple DAGs, this can work well by allowing the co-location of members from different DAGs on the same host. This allows you to drive consolidation higher, if the physical hardware supports the compute requirements.

Example

The flexibility of vSphere allows for multiple deployment options. In the previous examples, four DAG nodes were used for sizing. To understand how scaling out might change compute requirements, another sizing exercise was performed with six DAG nodes. The details of this scenario are as follows.

- Deploying six Exchange mailbox virtual machines requires each virtual machine to support 4,000 mailboxes during normal operation and up to 4,800 during a single server failure.
- Each mailbox virtual machine requires 6 vCPUs and 64GB of memory.
- Each mailbox server virtual machine supports 12 databases.
- The vSphere cluster must provide the following:
  - At least 36 physical CPU cores.
  - 384GB of memory.

Scaling out provides for better agility, reduced total compute resources, and fewer databases per mailbox virtual machine to manage. However, more DAG members require more ESXi hosts to keep DAG members on separate physical hosts.

2.6 Sample Physical Layout

Using the initial sizing example of four DAG nodes, the physical layout of virtual machines is illustrated in the following figure. The spare capacity within the ESXi hosts is used to provide resources for client access servers and any other peripheral services used for Exchange, such as backup or archive systems.

Figure 1. Sample Physical Layout for 24,000 Mailboxes

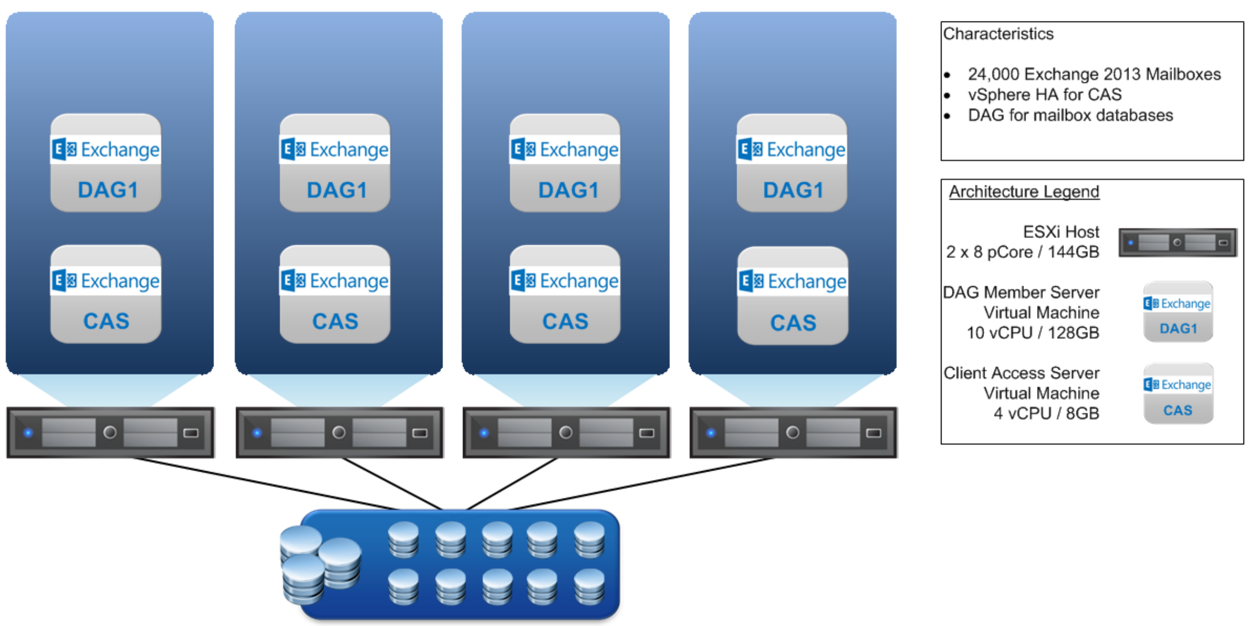

3. Sizing Examples

The following examples are provided to illustrate the topics covered in this guide across multiple deployment scenarios. These examples are meant to help reinforce the methodology for sizing an Exchange environment on vSphere and understand the flexibility available based on your deployment requirements and constraints.

Note Processor utilization, memory sizing, and I/O estimations are based on Exchange 2010 sizing guidance with considerations taken for Exchange 2013 architecture changes. Although sizing might change as Microsoft releases updated guidance for Exchange 2013, the methodology remains the same at the vSphere level.

In each of these examples, the following design parameters are used:

- Database size – Default.
- Average mailbox quota – 2048MB.
- Average messages sent and received per day – 150.
- Average message size – 75KB.
- Deleted item retention – 14 days.
- IOPS and megacycle multiplication factor – 1.00.
- Processor SPECint2006 rating – 8 cores per processor, 41 per core.

3.1 Single Role Server Design (12,000 Users)

This example uses separate Exchange virtual machines for both Client Access and Mailbox server roles. vSphere hosts are sized to provide failover capacity for all virtual machines. In this design two vSphere hosts can be taken offline with no impact to performance or further consolidation can be achieved.

3.1.1 Resource Requirements by Server Role

The following table lists the compute requirements for each server role to support 12,000 users.

Table 5. Exchange Server Role Resource Requirements

Exchange Role	Physical Resources per Server
Mailbox Server – 4 Servers	• CPU – 4 cores (31% max utilization). • Memory – 48GB. • OS and Application File Storage – 100GB (OS and application files). • Database Storage – 28 x 2000GB 7.2K RPM SAS 3.5" (RAID 1/0). • Log Storage – 2 x 2000GB 7.2K RPM SAS 3.5" (RAID 1/0). • Restore LUN – 3 x 2000GB 7.2K RPM SAS 3.5" (RAID 5). • Network – 1Gbps.
Client Access Server – 4 Servers	• CPU – 4 cores. • Memory – 8GB. • Storage – 80GB (OS and application files). • Network – 1Gbps.

3.1.2 Guest Virtual Machine Configuration

The resource requirements in the preceding table are translated into the following virtual machine resources.

Table 6. Exchange Virtual Machine Configuration

Exchange Role	Drive Letter/ Mount Point	Virtual Hardware per Virtual Machine
• Mailbox Server – 4 Servers • Normal Run Time – 3000 Mailboxes Each		• CPU – 4 vCPU. • Memory – 48GB. • Storage – SCSI Controller 1:
	C:\	○ HDD 1 – 80GB (OS and application files).
	D:\	○ HDD 2 – 1852GB (DB1 – DB7).
	E:\	○ HDD 3 – 83GB (LOG1 – LOG7).
	F:\	○ HDD 4 – 1852GB (DB8 – DB14).
	G:\	○ HDD 5 – 83GB (LOG8 – LOG14).
		• Storage – SCSI Controller 2:
	H:\	○ HDD 6 – 1852GB (DB15 – DB21).
	I:\	○ HDD 7 – 83GB (LOG15 – LOG21).
	J:\	○ HDD 8 – 1852GB (DB22 – DB28).
	K:\	○ HDD 9 – 83GB (LOG22 – LOG28).
		• Storage – SCSI Controller 3:
	L:\	○ HDD 10 – 1852GB (DB29 – DB35).
	M:\	○ HDD 11 – 83GB (LOG29 – LOG35).
	N:\	○ HDD 12 – 1852GB (DB36 – DB42).
	O:\	○ HDD 13 – 83GB (LOG36 – LOG42).
		• Network – vNIC 1 – LAN/Client Connectivity.
Client Access Server – 4 Servers		• CPU – 2 vCPUs. • Memory – 8GB. • Storage – SCSI Controller 1: ○ 80GB (OS and application files). • Network – vNIC 1 – LAN/Client Connectivity.

3.1.3 Guest Virtual Machine Storage Interaction

The following figure illustrates the mailbox virtual machine configuration.

Figure 2. Mailbox Virtual Machine Configuration

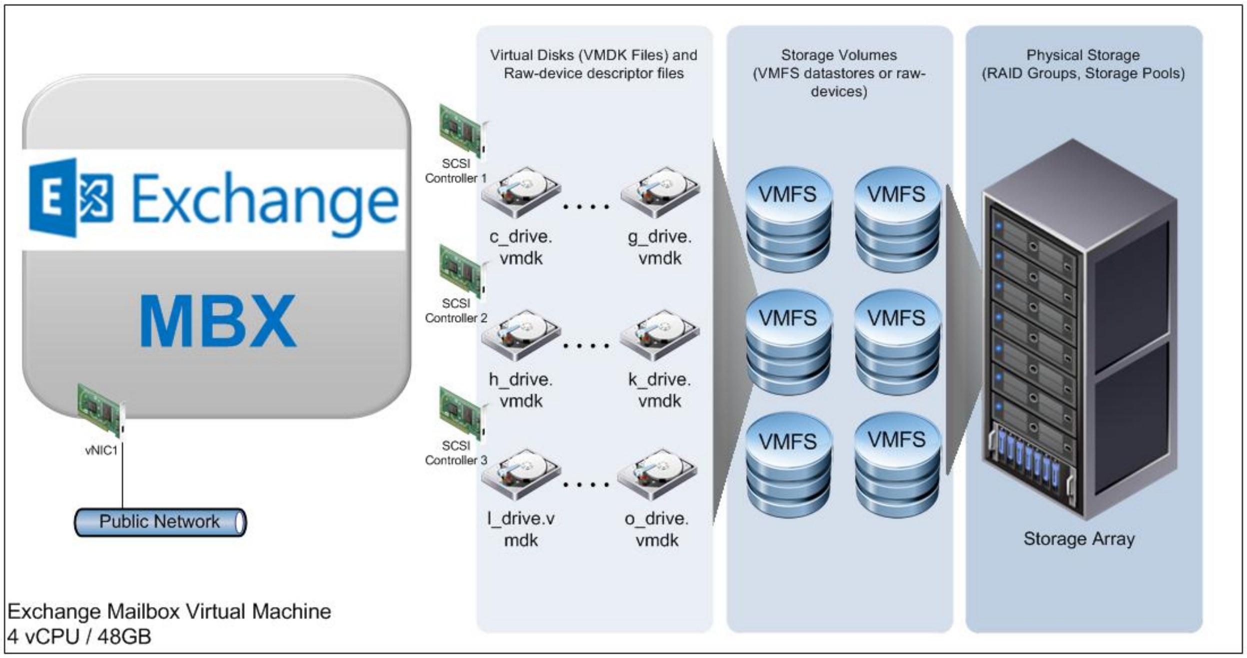

3.1.4 Virtual Machine Distribution

With an understanding of the physical resource requirements and associated virtual hardware configuration, you can plan for physical ESXi host hardware to meet those requirements. To build infrastructure availability into the architecture, distribute the eight total virtual machines across four ESXi hosts. Initial placement of virtual machines is relatively unimportant, especially if you are using DRS.

Table 7. Exchange Virtual Machine Distribution

ESXi Host	Virtual Machines
ESXi Host 1	• Exchange Mailbox Virtual Machine 1 – 4 vCPU, 48GB RAM. • Exchange Client Access Virtual Machine 1 – 2 vCPU, 8GB RAM.
ESXi Host 2	• Exchange Mailbox Virtual Machine 2 – 4 vCPU, 48GB RAM. • Exchange Client Access Virtual Machine 2 – 2 vCPU, 8GB RAM.
ESXi Host 3	• Exchange Mailbox Virtual Machine 3 – 4 vCPU, 48GB RAM. • Exchange Client Access Virtual Machine 3 – 2 vCPU, 8GB RAM.
ESXi Host 4	• Exchange Mailbox Virtual Machine 4 – 4 vCPU, 48GB RAM. • Exchange Client Access Virtual Machine 4 – 2 vCPU, 8GB RAM.

3.1.5 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 vSphere HA failover or planned vSphere vMotion migration of live virtual machines for host hardware maintenance. The following table summarizes the ESXi host hardware configuration based on the example architecture.

Table 8. ESXi Host Hardware Configuration Table

ESXi Host	Virtual Machines
All ESXi Hosts	• 16 cores – 2x 8. • 128GB RAM. • 2x Fibre Channel HBAs. • 4x Gigabit network adapters.

3.1.6 Initial Virtual Machine Placement

Although all of the virtual machines can migrate automatically with DRS, the following diagram is a useful planning tool for initial placement of virtual machines and for calculating host failover capacity. At initial placement, all ESXi hosts have some failover headroom. During a single host failure, vSphere HA can power on any failed virtual machines, bringing them back into service.

Figure 3. Initial Virtual Machine Placement

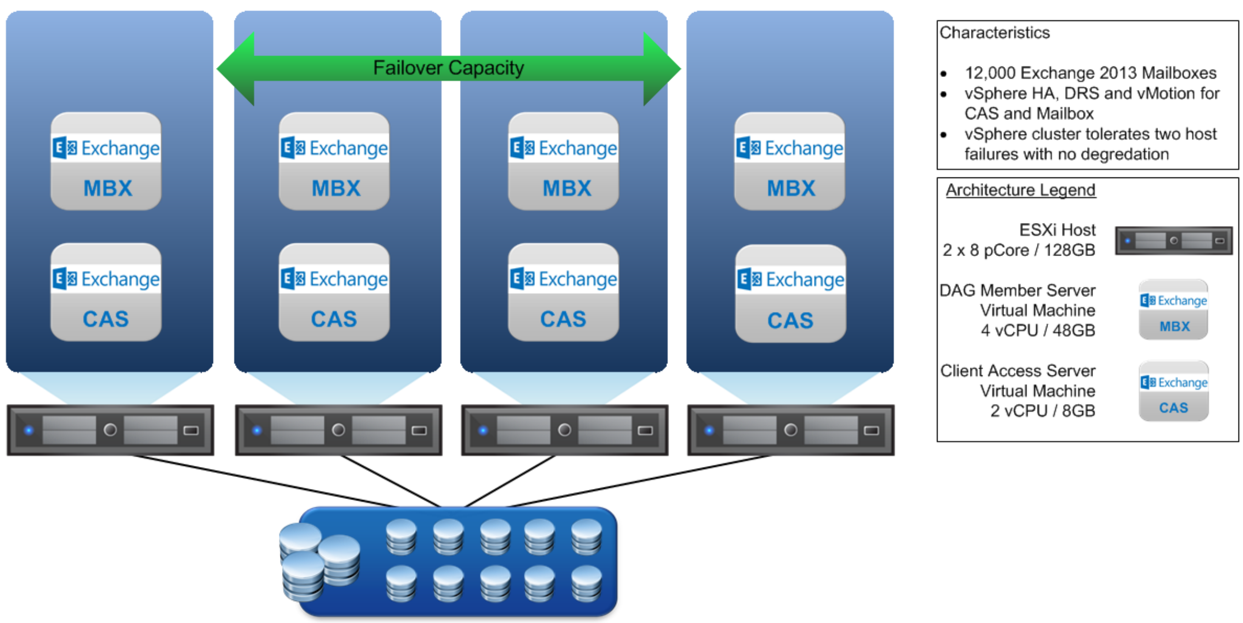

3.2 Single Role Server Design with DAG for 24,000 Users

This example uses dedicated Exchange virtual machines for both the Client Access and Mailbox server roles. Client Access servers are protected using vSphere HA. Mailbox servers are protected by a DAG. vSphere hosts are sized to provide some failover capacity for client access servers but could also be used to house other peripheral service virtual machines.

3.2.1 Resource Requirements by Server Role

The following table lists the compute requirements for each server role to support 24,000 users.

Table 9. Exchange Server Role Resource Requirements

Exchange Role	Physical Resources per Server
Mailbox Server – 4 Servers	• CPU –10 cores (39% max utilization). • Memory – 128GB. • OS and Application File Storage – 100GB (OS and application files). • Database Storage – 82x 2000GB 7.2K RPM SAS 3.5" (RAID 1/0). • Log Storage – 6x 2000GB 7.2K RPM SAS 3.5". • Restore LUN – 12x 2000GB 7.2K RPM SAS 3.5" (RAID 5). • Network –1Gbps.
Client Access Server – 4 Servers	• CPU – 4 cores. • Memory – 8GB. • Storage – 80GB (OS and application files). • Network – 1Gbps.

3.2.2 Guest Virtual Machine Configuration

The resource requirements in the preceding table are translated into the following virtual machine resources.

Table 10. Exchange Virtual Machine Configuration

Exchange Role	Drive Letter/ Mount Point	Virtual Hardware per Virtual Machine

Mailbox Server – 3 Servers
Normal Run Time –6000 Mailboxes Each
After Single Server Failure – 8000 Mailboxes Each
	C:\
	D:\
	E:\
	F:\
	G:\
	H:\
	I:\
	J:\
		• CPU – 10 vCPU.
		• Memory – 128GB.
		• Storage – SCSI Controller1:
		  ○ HDD 1 – 80GB (OS and application files).
		  ○ HDD 2 – 14818GB (DB1 – DB6).
		  ○ HDD 3 – 662GB (LOG1 – LOG6).
		• Storage – SCSI Controller 2:
		  ○ HDD 4 – 14818GB (DB7 – DB12).
		  ○ HDD 5 – 662GB (LOG7 – LOG12).
		• Storage – SCSI Controller 3:
		  ○ HDD 6 – 14818GB (DB13 – DB18).
		  ○ HDD 7 – 662GB (LOG13 – LOG18).
		• Storage: SCSI Controller 4:
		  ○ HDD 8 – 14133GB (Restore LUN).
		• Network:
		  ○ vNIC 1 – LAN/Client Connectivity.
		  ○ vNIC 2 – DAG Replication/Heartbeat.

Client Access Server – 4 Servers		• CPU – 4 vCPU. • Memory – 8GB. • Storage – SCSI Controller 1: ○ 80GB (OS and application files). • Network – vNIC 1 – LAN/Client Connectivity.

3.2.3 Guest Virtual Machine Storage Interaction

The following figure illustrates the per virtual machine configuration.

Figure 4. Mailbox Virtual Machine Configuration

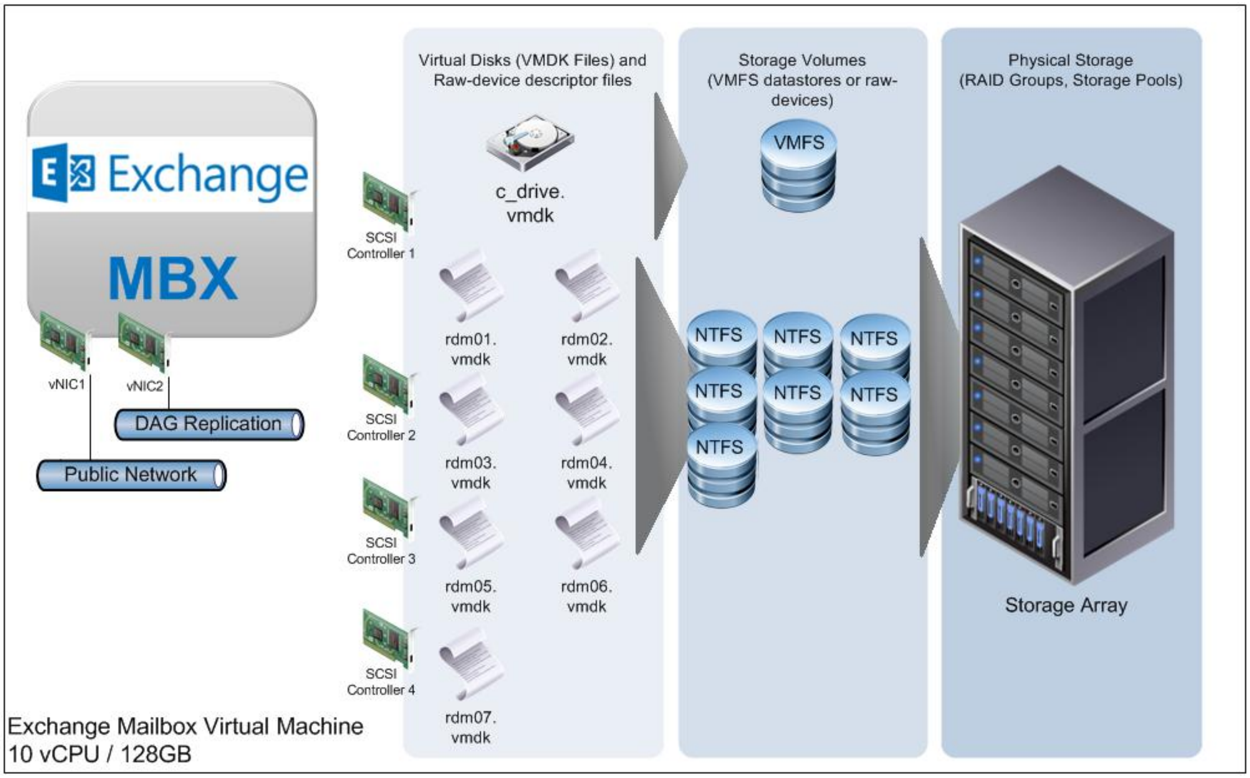

3.2.4 Virtual Machine Distribution

After understanding the physical resource requirements and associated virtual hardware configuration, you can plan physical ESXi host hardware to meet those requirements. To build infrastructure availability into the architecture, distribute the eight total virtual machines across four ESXi hosts. Initial placement of virtual machines relies on DRS anti-affinity rules to keep DAG members separated.

Table 11. Exchange Virtual Machine Distribution

ESXi Host	Virtual Machines
ESXi Host 1	• Exchange Mailbox Virtual Machine 1 – 10 vCPU, 128GB RAM. • Exchange Client Access Virtual Machine 1 – 4 vCPU, 8GB RAM.
ESXi Host 2	• Exchange Mailbox Virtual Machine 2 – 10 vCPU, 128GB RAM. • Exchange Client Access Virtual Machine 2 – 4 vCPU, 8GB RAM.
ESXi Host 3	• Exchange Mailbox Virtual Machine 3 – 10 vCPU, 128GB RAM. • Exchange Client Access Virtual Machine 3 – 4 vCPU, 8GB RAM.
ESXi Host 4	• Exchange Mailbox Virtual Machine 4 – 10 vCPU, 128GB RAM. • Exchange Client Access Virtual Machine 4 – 4 vCPU, 8GB RAM.

3.2.5 ESXi Host Specifications

Each ESXi host should provide enough physical hardware resources to accommodate the planned workload. The following table summarizes the ESXi host hardware configuration based on our example architecture.

Table 12. ESXi Host Hardware Configuration Table

ESXi Host	Virtual Machines
All ESXi Hosts	• 16 cores – 2x 8. • 144GB RAM. • 2s Fibre Channel HBAs. • 4s Gigabit network adapters.

3.2.6 Initial Virtual Machine Placement

The following diagram is a useful planning tool for initial placement of virtual machines and for calculating host capacity. During a single host failure, vSphere HA can power on the failed Client Access server virtual machine on a remaining ESXi host. The mailbox virtual machine stays powered off, and the DAG activates any affected databases. All mailboxes continue to be accessible. To provide failover capabilities for all virtual machines, including mailbox virtual machines, a fifth ESXi host can be added into the vSphere cluster. This allows the DAG to be brought back into a protected state automatically.

Figure 5. Initial Virtual Machine Placement

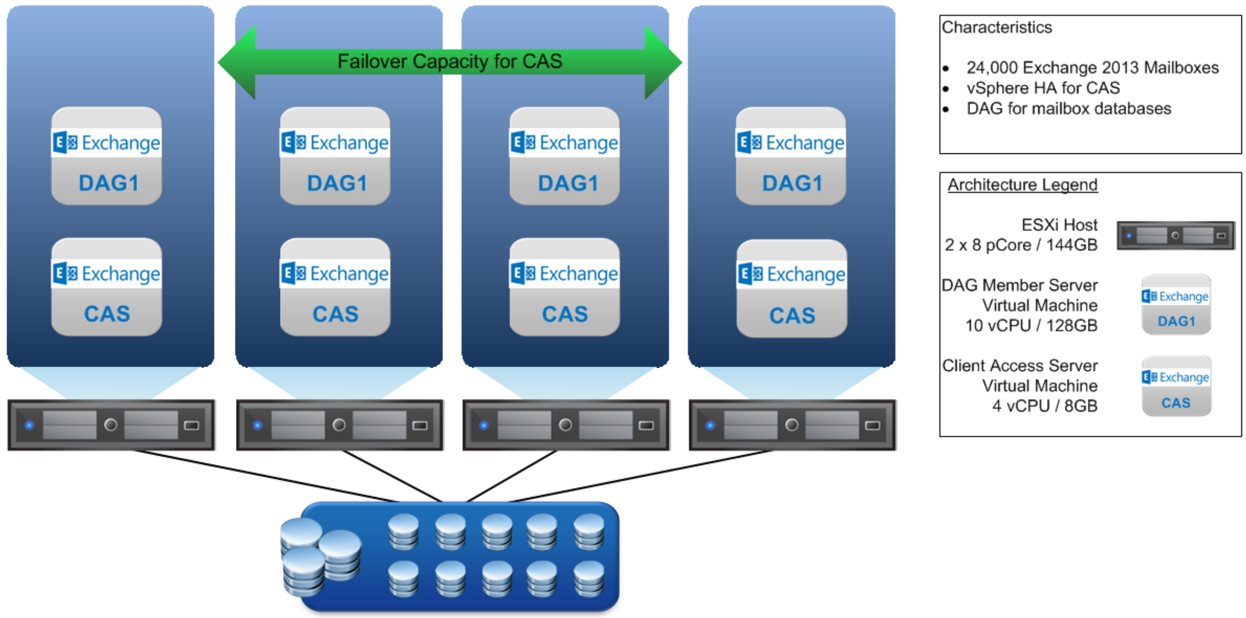

3.3 Multirole and Multisite DAG Server Design (50,000 Users)

This example uses multirole Exchange virtual machines with the Client Access and Mailbox server role installed in a single virtual machine. All virtual machines are protected using vSphere HA and DAG. Site resilience is provided using DAG to replicate data to a third database copy in the secondary datacenter. Two DAGs are used to leverage the full capabilities of the ESXi host resources and to be able to run multiple mailbox virtual machines on a single host. A physical deployment requires additional hosts to provide site resiliency. With vSphere, the available compute resources can be used.

3.3.1 Resource Requirements by Server Role

The following table lists the compute requirements to support 50,000 users across two DAGs.

Table 13. Exchange Server Role Resource Requirements

Exchange Role	Physical Resources per Server
Primary Datacenter Mailbox Server – 12 Servers	• CPU – 8 cores (34% max utilization). • Memory – 96GB. • OS and Application File Storage –100GB (OS and application files). • Database Storage – 58x 2000GB 7.2K RPM SAS 3.5" (RAID 1/0). • Log Storage – 4x 2000GB 7.2K RPM SAS 3.5". • Restore LUN – 3x 2000GB 7.2K RPM SAS 3.5" (RAID 5). • Network – 1Gbps.
Secondary Datacenter Mailbox Server – 6 Servers	• CPU – 8 cores (51% max utilization). • Memory – 128GB. • OS and Application File Storage – 100GB (OS and application files). • Database Storage – 58x 2000GB 7.2K RPM SAS 3.5" (RAID 1/0). • Log Storage – 4x 2000GB 7.2K RPM SAS 3.5". • Restore LUN – 3x 2000GB 7.2K RPM SAS 3.5" (RAID 5). • Network –1Gbps.

3.3.2 Guest Virtual Machine Configuration

The resource requirements in the preceding table are translated into the following virtual machine resources.

Table 14. Exchange Virtual Machine Configuration

Exchange Role	Drive Letter/ Mount Point	Virtual Hardware per Virtual Machine
• Primary Datacenter Mailbox Server – 12 servers • Datacenter A – 6 Servers • Datacenter B – 6 Servers • Normal Run Time – 4167 Mailboxes Each • After Single Server Failure – 5060 Mailboxes Each	 C:\ C:\Mounts\DB1 C:\Mounts\LOG1 C:\Mounts\DB2 C:\Mounts\LOG2 C:\Mounts\DB3 C:\Mounts\LOG3 C:\Mounts\DB4 C:\Mounts\LOG4 C:\Mounts\DB5 C:\Mounts\LOG5 C:\Mounts\DB6 C:\Mounts\LOG6 C:\Mounts\DB7 C:\Mounts\LOG7 C:\Mounts\DB8 C:\Mounts\LOG8 C:\Mounts\DB9 C:\Mounts\LOG9 C:\Mounts\DB10 C:\Mounts\LOG10 C:\Mounts\DB11 C:\Mounts\LOG11 C:\Mounts\DB12 C:\Mounts\LOG12 C:\Mounts\Restore	• CPU – 8 vCPU • Memory – 96GB • Storage – SCSI Controller 1 ○ HDD 1 – 80GB (OS and application files) ○ HDD 2 – 2573GB (DB1) ○ HDD 3 – 115GB (LOG1) ○ HDD 4 – 2573GB (DB2) ○ HDD 5 – 115GB (LOG2) ○ HDD 6 – 2573GB (DB3) ○ HDD 7 – 115GB (LOG3) • Storage – SCSI Controller 2 ○ HDD 8 – 2573GB (DB4) ○ HDD 9 – 115GB (LOG4) ○ HDD 10 – 2573GB (DB5) ○ HDD 11 – 115GB (LOG5) ○ HDD 12 – 2573GB (DB6) ○ HDD 13 – 115GB (LOG6) • Storage – SCSI Controller 3 ○ HDD 14 – 2573GB (DB7) ○ HDD 15 – 115GB (LOG7) ○ HDD 16 – 2573GB (DB8) ○ HDD 17 – 115GB (LOG8) ○ HDD 18 – 2573GB (DB9) ○ HDD 19 – 115GB (LOG9) • Storage – SCSI Controller 4 ○ HDD 20 – 2573GB (DB10) ○ HDD 21 – 115GB (LOG10) ○ HDD 22 – 2573GB (DB11) ○ HDD 23 – 115GB (LOG11) ○ HDD 24 – 2573GB (DB12) ○ HDD 25 – 115GB (LOG12) ○ HDD 26 – 2454GB (Restore LUN) • Network ○ vNIC 1 – LAN/Client Connectivity ○ vNIC 2 – DAG Replication/Heartbeat

Exchange Role	Drive Letter/ Mount Point	Virtual Hardware per Virtual Machine
• Secondary Datacenter Mailbox Server – 6 Servers		• CPU – 8 vCPU
• Datacenter A – 3 Servers		• Memory – 128GB
• Datacenter B – 3 Servers		• Storage – SCSI Controller 1
• If Fully Activated – 8333 Mailboxes Each	C:\	○ HDD 1 – 80GB (OS and application files)
	C:\Mounts\DB1	○ HDD 2 – 2573GB (DB1)
	C:\Mounts\LOG1	○ HDD 3 – 115GB (LOG1)
	C:\Mounts\DB2	○ HDD 4 – 2573GB (DB2)
	C:\Mounts\LOG2	○ HDD 5 – 115GB (LOG2)
	C:\Mounts\DB3	○ HDD 6 – 2573GB (DB3)
	C:\Mounts\LOG3	○ HDD 7 – 115GB (LOG3)
		• Storage – SCSI Controller 2
	C:\Mounts\DB4	○ HDD 8 – 2573GB (DB4)
	C:\Mounts\LOG4	○ HDD 9 – 115GB (LOG4)
	C:\Mounts\DB5	○ HDD 10 – 2573GB (DB5)
	C:\Mounts\LOG5	○ HDD 11 – 115GB (LOG5)
	C:\Mounts\DB6	○ HDD 12 – 2573GB (DB6)
	C:\Mounts\LOG6	○ HDD 13 – 115GB (LOG6)
		• Storage – SCSI Controller 3
	C:\Mounts\DB7	○ HDD 14 – 2573GB (DB7)
	C:\Mounts\LOG7	○ HDD 15 – 115GB (LOG7)
	C:\Mounts\DB8	○ HDD 16 – 2573GB (DB8)
	C:\Mounts\LOG8	○ HDD 17 – 115GB (LOG8)
	C:\Mounts\DB9	○ HDD 18 – 2573GB (DB9)
	C:\Mounts\LOG9	○ HDD 19 – 115GB (LOG9)
		• Storage – SCSI Controller 4
	C:\Mounts\DB10	○ HDD 20 – 2573GB (DB10)
	C:\Mounts\LOG10	○ HDD 21 – 115GB (LOG10)
	C:\Mounts\DB11	○ HDD 22 – 2573GB (DB11)
	C:\Mounts\LOG11	○ HDD 23 – 115GB (LOG11)
	C:\Mounts\DB12	○ HDD 24 – 2573GB (DB12)
	C:\Mounts\LOG12	○ HDD 25 – 115GB (LOG12)
	C:\Mounts\Restore	○ HDD 26 – 2454GB (Restore LUN)
		• Network
		○ vNIC 1 – LAN/Client Connectivity
		○ vNIC 2 – DAG Replication/Heartbeat

3.3.3 Guest Virtual Machine Storage Interaction

The following figure illustrates the per virtual machine configuration.

Figure 6. Mailbox Virtual Machine Configuration

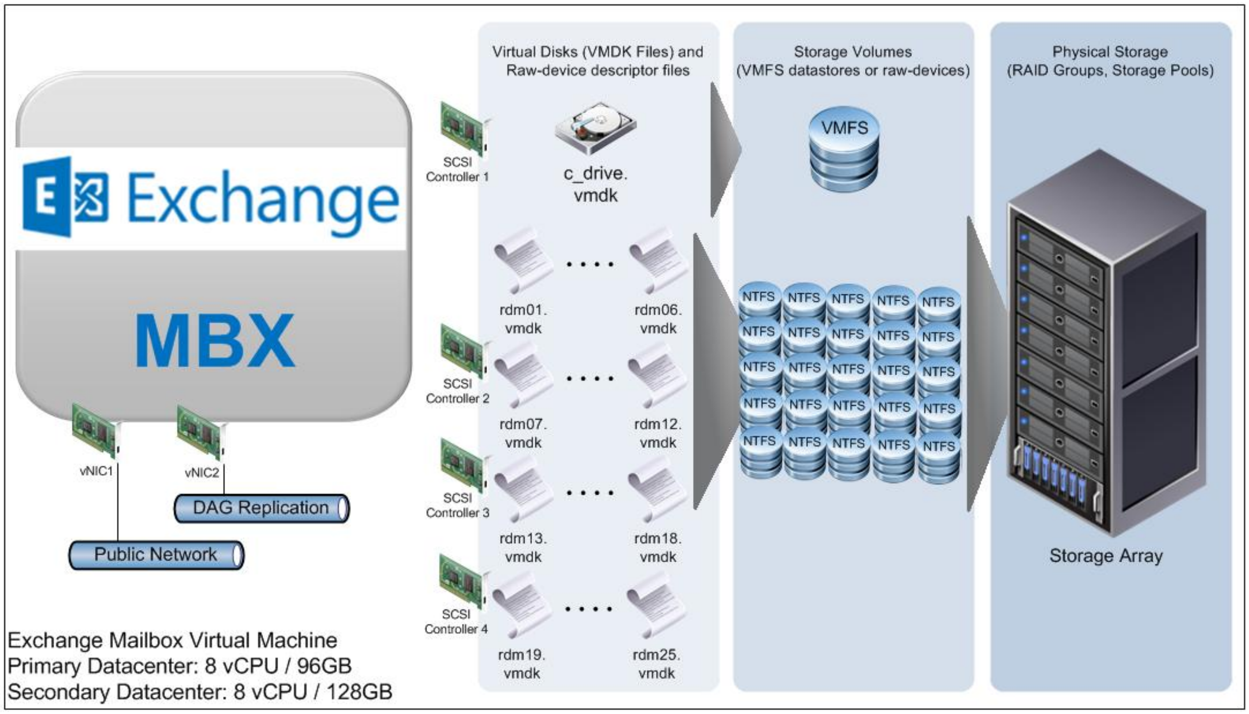

3.3.4 Virtual Machine Distribution

With an understanding of the physical resource requirements and the associated virtual hardware configuration, you can plan physical ESXi host hardware to meet those requirements. To build infrastructure availability into the architecture, distribute the nine virtual machines across six ESXi hosts in each site. Initial placement of virtual machines relies on DRS anti-affinity rules to keep DAG members separated.

Table 15. Exchange Virtual Machine Distribution: Datacenter A

ESXi Host	Virtual Machines
ESXi Host 1	• DAG1-Exchange Mailbox Virtual Machine 1 – 8 vCPU, 96GB RAM. • DAG2-Exchange Mailbox Virtual Machine 7 – 8 vCPU, 128GB RAM.
ESXi Host 2	• DAG1-Exchange Mailbox Virtual Machine 2 – 8 vCPU, 96GB RAM. • DAG2-Exchange Mailbox Virtual Machine 8 – 8 vCPU, 128GB RAM.
ESXi Host 3	• DAG1-Exchange Mailbox Virtual Machine 3 – 8 vCPU, 96GB RAM. • DAG2-Exchange Mailbox Virtual Machine 9 – 8 vCPU, 128GB RAM.
ESXi Host 4	• DAG1-Exchange Mailbox Virtual Machine 4 – 8 vCPU, 96GB RAM.

ESXi Host	Virtual Machines
ESXi Host 5	• DAG1-Exchange Mailbox Virtual Machine 5 – 8 vCPU, 96GB RAM.
ESXi Host 6	• DAG1-Exchange Mailbox Virtual Machine 6 – 8 vCPU, 96GB RAM.

Table 16. Exchange Virtual Machine Distribution: Datacenter B

ESXi Host	Virtual Machines
ESXi Host 7	• DAG2-Exchange Mailbox Virtual Machine 1 – 8 vCPU, 96GB RAM. • DAG1-Exchange Mailbox Virtual Machine 7 – 8 vCPU, 128GB RAM.
ESXi Host 8	• DAG2-Exchange Mailbox Virtual Machine 2 – 8 vCPU, 96GB RAM. • DAG1-Exchange Mailbox Virtual Machine 8 – 8 vCPU, 128GB RAM.
ESXi Host 9	• DAG2-Exchange Mailbox Virtual Machine 3 – 8 vCPU, 96GB RAM. • DAG1-Exchange Mailbox Virtual Machine 9 – 8 vCPU, 128GB RAM.
ESXi Host 10	• DAG2-Exchange Mailbox Virtual Machine 4 – 8 vCPU, 96GB RAM.
ESXi Host 11	• DAG2-Exchange Mailbox Virtual Machine 5 – 8 vCPU, 96GB RAM.
ESXi Host 12	• DAG2-Exchange Mailbox Virtual Machine 6 – 8 vCPU, 96GB RAM.

3.3.5 ESXi Host Specifications

Each ESXi host should provide enough physical hardware resources to accommodate the planned workload. The following table summarizes the ESXi host hardware configuration based on the example architecture.

Table 17. ESXi Host Hardware Configuration Table

ESXi Host	Virtual Machines
All ESXi Hosts	• 16 cores – 2x 8. • 256GB RAM. • 2x Fibre Channel HBAs. • 4x Gigabit network adapters.

3.3.6 Initial Virtual Machine Placement

The following diagram is a useful planning tool for initial placement of virtual machines and for calculating host capacity. During a single host failure, vSphere HA can power on the failed secondary datacenter mailbox virtual machine on a remaining ESXi host. The primary datacenter mailbox virtual machine stays powered off, and the DAG activates any affected databases. All mailboxes continue to be accessible. To provide for failover capabilities for all virtual machines, including mailbox virtual machines, a seventh ESXi host can be added into the vSphere cluster in each site. This allows the DAG to be brought back into a protected state automatically.

Figure 7. Initial Virtual Machine Placement

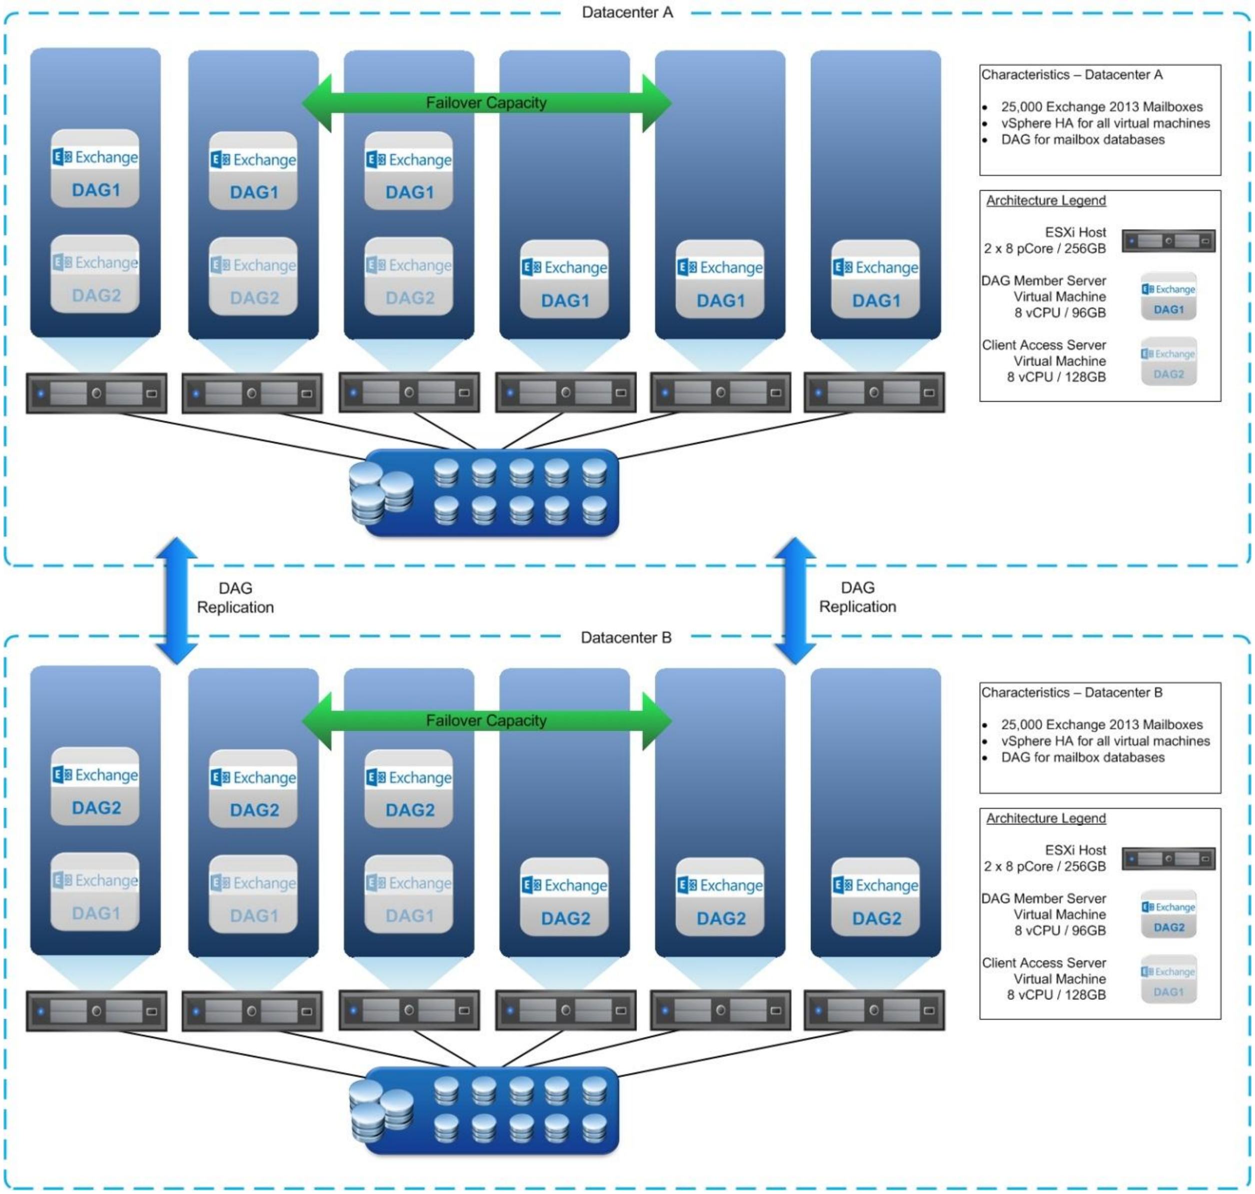

4. Summary

Much of the design and sizing considerations that formerly went into a physical Exchange design continue to apply in a virtualized environment. Functional elements of the Exchange design, such as name space design and mail flow, remain unchanged due to virtualization. Other design considerations, such as high availability, disaster recovery, and server sizing, can use functions of vSphere, Exchange, or a hybrid of both to provide the best possible environment for Exchange 2013.

Knowledge of vSphere capabilities and configuration maximums help to verify that the Exchange design fits into a vSphere architecture. The flexibility provided by vSphere allows Exchange architects to size Exchange servers to meet the design requirements while providing efficient use of physical compute resources.

This guide shows example configurations of Exchange 2013 on VMware. These examples provide high-level guidance and are not intended to reflect customer-specific workloads. Customers should work with infrastructure vendors to build a detailed sizing and architecture design that meets individual requirements.