

Chapter 6

Virtualizing Exchange Server

Messaging and collaboration have become critical components in most if not all business operations. In businesses, email is integrated into most business processes, including the help desk, bug tracking, archiving, document management, and monitoring operations. When email is unavailable, business productivity decreases, and foot traffic in the break room increases dramatically. Escalations and calls to the help desk increase dramatically. There are many reasons for an Exchange server outage to include server hardware failure, database corruption, and performance degradation because of server storage over-subscription. Virtualizing Exchange Server can help mitigate most of the issues associated with a physical deployment by doing the following:

- ◆ Allowing you to right-size your Exchange Server deployment by adding resources as needed with just a reboot.
- ◆ Moving your Exchange Server to a more capable host to perform host upgrades without downtime to your end users through the use of VMotion.
- ◆ Providing simplified availability through the use of VMware HA if your Exchange Server virtual machines hang or you get a blue screen.
- ◆ Allowing you to create simplified disaster recovery scenarios without the requirement of similar hardware. A virtual machine can run on any VMware-supported physical server that has the ESX hypervisor installed.

Over the past 10 years, the Microsoft Exchange Server product team has managed to create a messaging and collaboration product that captures more than 50 percent of the messaging and collaboration market share all while dominating this space.

Microsoft has worked very hard on this release to make the product easier to design, deploy, manage, and recover. With all the new features and enhancements, Microsoft Exchange 2010 is threatening to dominate even more of the messaging and collaboration market. There are a lot of compelling reasons to transition from an Exchange Server 2003 or 2007 environment to an Exchange 2010 environment.

Exchange Server 2010 is totally virtualization friendly, more so than Exchange 2003 and Exchange 2007. We will briefly cover some of the more important features and reasons you might want to start planning for a virtualized Exchange 2007 or 2010 environment.

Exchange Server Virtualization Trends

Prior to 2007, not many medium to large organizations were even thinking about virtualizing Exchange Server, let alone actually doing it. These were the major objections to virtualizing this application:

- ◆ Lack of Microsoft support for the virtualized Exchange Server application
- ◆ Virtualization overhead
- ◆ Poor application performance when virtualized
- ◆ Poor storage I/O performance when virtualized
- ◆ The “If it ain’t broke, don’t fix it” mentality
- ◆ Fear of loss of employment

Some experimented with virtualizing Exchange Server 2003 but abandoned the effort when they experienced severe performance issues with Mailbox server virtualization.

Microsoft also didn’t make it easy to virtualize its applications, which created some major barriers to the virtualization of tier-1 applications. In the fall of 2008, Microsoft announced the Server Virtualization Validation Program (SVVP) and, in doing so, took a major step forward into the virtualization journey.

With this program, Microsoft announced support for its major applications, and Exchange was one of the first. Before this announcement of support, some brave and daring customers took a modular approach and virtualized the bridgehead servers and front-end servers and left the back-end Mailbox servers on physical servers.

Despite the warnings from Microsoft that they would not be supported or that they would be required to reproduce the issue on physical hardware in order to receive support, customers continued to pursue virtualization.

With the maturity of the VMware virtualization platforms and the formation of the tier-1 applications performance team within VMware, VMware was able to extensively test Exchange Server 2003 and Exchange 2007 and produce recommendations and guidance for virtualizing the Exchange products. After producing these recommendations and guidance, VMware started to see an uptick in the virtualization of the Exchange Server product. Most of this virtualization was being done on Exchange Server 2007 with minimal efforts being applied to the Exchange Server 2003 platform.

VMware then started to work with the engineering folks to enlist the various server and storage manufacturing partners to improve the performance of virtualized tier-1 applications. Microsoft Exchange Server was one of the first tier-1 applications with which VMware did extensive performance testing and performance optimizations on the VMware virtualization platform.

To prove to the information technology world that the virtualization of Exchange was entirely possible, VMware took on the challenge of virtualizing its internal Exchange environment.

Two authors of this book, Alex Fontana and Charles Windom, were the principal architects for the design and implementation. During this process, the architects worked directly with VMware engineers, members of the file system group, and the high-availability folks within VMware to ensure that the virtualization project would meet or exceed the performance of the existing Exchange Server 2003 implementation.

TIP For more information on the VMware Internal Exchange virtualization design and implementation, please visit the VMware website at www.vmware.com/solutions/business-critical-apps/exchange/resources.html.

In an effort to improve the performance for the virtualized Exchange deployments, VMware started working with its server and storage partners to create real-world Exchange 2003 and 2007 configurations. Today, work is still being done to constantly improve the performance of the vSphere platform as it pertains to running the Microsoft 2010 suite of tier-1 applications.

Exchange Server Comparisons

As the Microsoft Exchange Server platform has continued to evolve over the years, it has changed from the typical all-in-one message server to a more distributed environment.

Exchange Server 2003

Exchange Server 2003 offered the following distributed server roles; some of these server roles were more difficult if not impossible to virtualize:

Front-end server This is responsible for external client connectivity (Outlook Web Access, POP3, IMAP, RPC/HTTP, SMTP, and Outlook Mobile Access, to name a few). This server role was easy to virtualize because it wasn't very demanding on storage resources, and the VMware platform provided more than adequate CPU, memory, and network resources for these roles to perform as well as a physical server.

Exchange bridgehead server This is responsible for delivering messages to Mailbox servers outside the client sender site. Exchange Server 2003 utilizes routing groups to send and receive messages between Exchange sites, such as a local Exchange site in Palo Alto and a remote site in Japan. This was another server role that could be virtualized, although you needed to pay more attention to the configuration of the virtual machine for this role because it required more CPU, memory, network, and storage resources than the Exchange front-end server role.

Back-end server This server role contains the database(s) used to hold the mail databases. This is also the server role that provides mailbox and public folder access. It can also (optionally) provide NNTP, POP, and IMAP4 access. This server role required careful design and planning when virtualizing.

Virtualizing this server role in large to very large organizations proved daunting if not impossible. Because of the demanding CPU and storage requirements for this role, these organizations would have to deploy more virtual machines and increase the management overhead associated with managing those virtual machines. This proved to be a major obstacle in virtualizing this server role.

Exchange Server 2007 and 2010 Server Roles

With the introduction of Exchange 2007, Microsoft changed the configuration and naming convention of the server roles. Instead of the typical three server roles, Microsoft introduced five new server roles; three of the five had somewhat identical functions to the previous server roles.

There are no longer the Exchange front-end, bridgehead, and back-end server roles. The new names of the roles going forward are as follows:

- ◆ Exchange Edge server role (EDGE)
- ◆ Exchange Client Access server role (CAS)
- ◆ Exchange Hub Transport server role (HUB)
- ◆ Exchange Mailbox server role (MBX)
- ◆ Exchange Unified Messaging server role (UM)

Each Exchange server role performs a specific function in the Exchange organization. For each role, only the Exchange server services get installed and started for that role. You no longer have to worry about disabling multiple Exchange services and deleting databases on server roles that do not require them.

In the next sections, we will give a brief explanation for each role's responsibility in the Exchange organization.

Edge Server Role

This server must be deployed at the perimeter of the network. It should not be a member of the domain.

Figure 6.1 highlights the Exchange 2007/2010 Edge server role. As you can see in the figure, there are two Edge servers deployed in the demilitarized zone (DMZ) of the network.

You can deploy multiple Edge servers in the perimeter network. When you do deploy multiple Edge servers, they will provide redundancy and failover capabilities for inbound email.

You also have the ability to load balance your SMTP traffic to multiple Edge servers by defining multiple mail exchanger (MX) records with the same priority in the DNS database for your mail domain.

The Edge server also provides messaging hygiene known as antispam and antivirus control for email incoming to your Exchange organization and outgoing to the Internet.

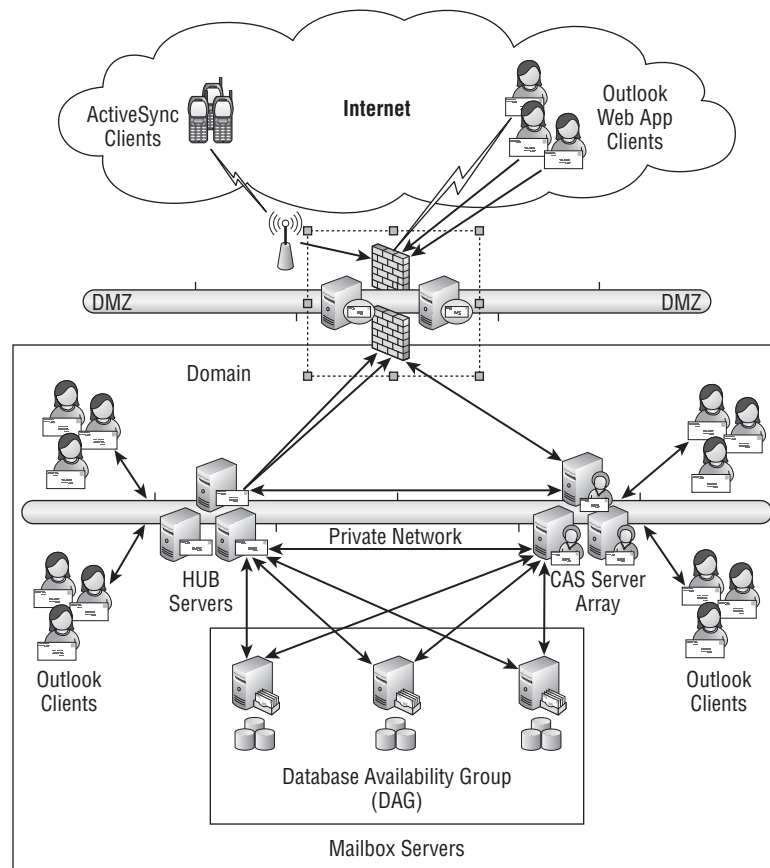
Edge transport servers can also prevent information that could be confidential to your organization from being sent outside the organization through the use of edge transport rules.

You can create and assign rules to control message flow out of the Exchange organization, and if a message meets the criteria of the created rule, that message will be blocked from being sent outside the organization.

The Edge transport server can also perform inbound and outbound SMTP address rewriting on messages. This can be especially useful to provide a consistent appearance of SMTP addresses in email organizations that perform lots of acquisitions.

Virtualizing this role is easy because it is not very demanding on CPU, memory, and disk resources. However, be mindful of your network resources when deploying multiple servers because this places additional demand on the physical network adapters in the ESX host. Ensure that you provide adequate network resources and redundancy through the use of ESX network teaming.

FIGURE 6.1
Edge server role



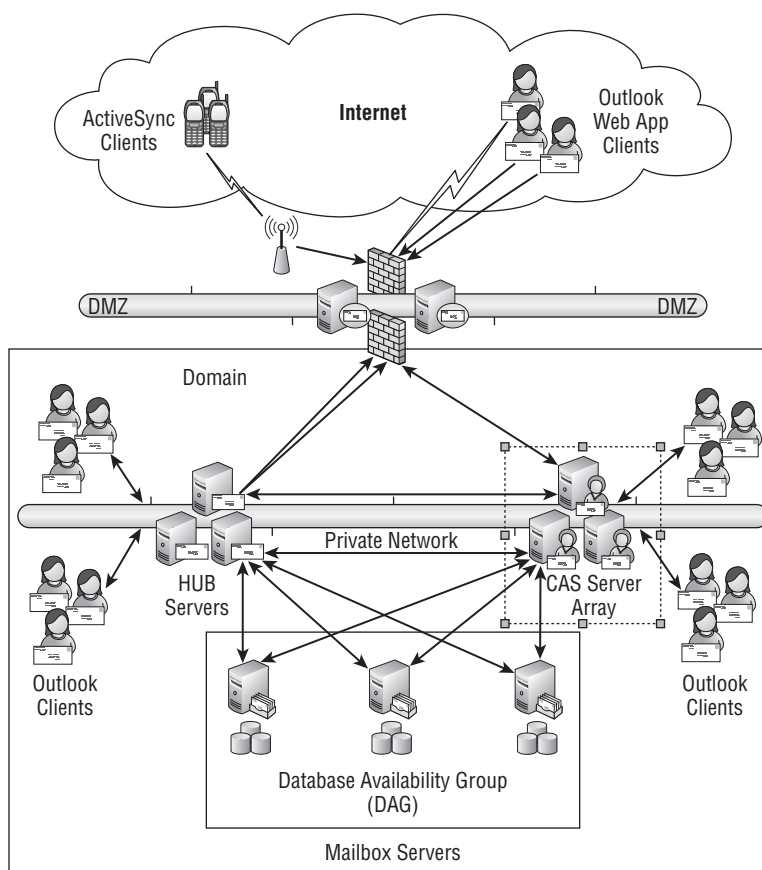
Client Access Server Role

The Client Access server (CAS) role hosts all client protocols such as Outlook Web Access (Outlook Web App for 2010), POP3, IMAP4, Outlook Mobile Access (OMA), and Outlook Anywhere (formerly known as RPC over HTTPS). The Client Access server role must be installed in every Active Directory site that has a Mailbox server installed.

Figure 6.2 shows the CAS role deployed as a CAS array on the internal network protected by the firewall.

The Client Access server is also responsible for providing access to the Exchange free/busy data and also allows Exchange 2007 and greater clients to download automatic configuration settings using the Autodiscover service.

FIGURE 6.2
Client Access
server array



To increase the scalability, availability, and redundancy of the Client Access server, in 2010 there were changes made to the Client Access server role. These are some of the fundamental changes:

- ◆ New client access server arrays (CAS arrays) provide availability and load balancing. The failure of a member of the CAS will automatically connect the client to another member of the CAS array.
- ◆ CAS arrays allow a higher number of concurrent connections per server and a higher number of mailboxes per server.
- ◆ The new RPC endpoint on the Client Access server role replaces the RPC endpoint on the Mailbox server role.
- ◆ New business logic routes Outlook clients to the CAS server to access a mailbox instead of logging onto the Mailbox server.

Exchange 2010 CAS server roles use RPC encryption by default, so Outlook 2007 and greater clients are compatible with Exchange 2010 CAS servers. Outlook 2003 clients will need to modify their settings to force RPC encryption in order to connect to a 2010 CAS server or array. Exchange Server or Windows Server administrators can also create a group policy to force Outlook 2003 to use RPC encryption.

This server role in Exchange 2007 can be easy to virtualize because you can also deploy multiple servers to provide load balancing and redundancy capabilities for the CAS server role. With Exchange Server 2010, be mindful that the CAS server role has taken on additional responsibilities (mentioned earlier in this section) and requires more server resources in the form of CPU, memory, disk, and network resources. The same rule of deploying multiple servers to provide load balancing and redundancy applies, but now you have the CAS array option. In actual production scenarios, we have seen that IMAP places additional demand on physical server resources than MAPI and could mean that you will have to deploy additional virtual machines to meet this additional demand.

TIP If your Outlook 2003 users cannot connect to an Exchange 2010 CAS server, or receive the error messages “Cannot start Microsoft Office Outlook. Unable to open the Office Window. The set of folders cannot be opened or unable to open your default email folders. The information store cannot be opened,” then try enabling the Encrypt Data Between Microsoft Office Outlook And Microsoft Exchange Server option in the Outlook 2003 security settings.

Hub Transport Role

The Hub Transport role is responsible for moving messages through the Exchange organization. When we say “through the Exchange organization,” we mean inside and outside the Exchange organization.

The Hub Transport role uses Active Directory sites to route email messages to local and remote Exchange sites in the email organization. Deploying multiple Hub transport servers in the same Active Directory site will provide availability as well as load balancing. There must be at least one Hub Transport server installed in every Active Directory site that has the Mailbox server role installed.

In the absence of the Edge server role, the Hub Transport can be configured to send and receive messages to and from the Internet as well as provide messaging hygiene, messaging policy, and compliance using compliance and transport rules.

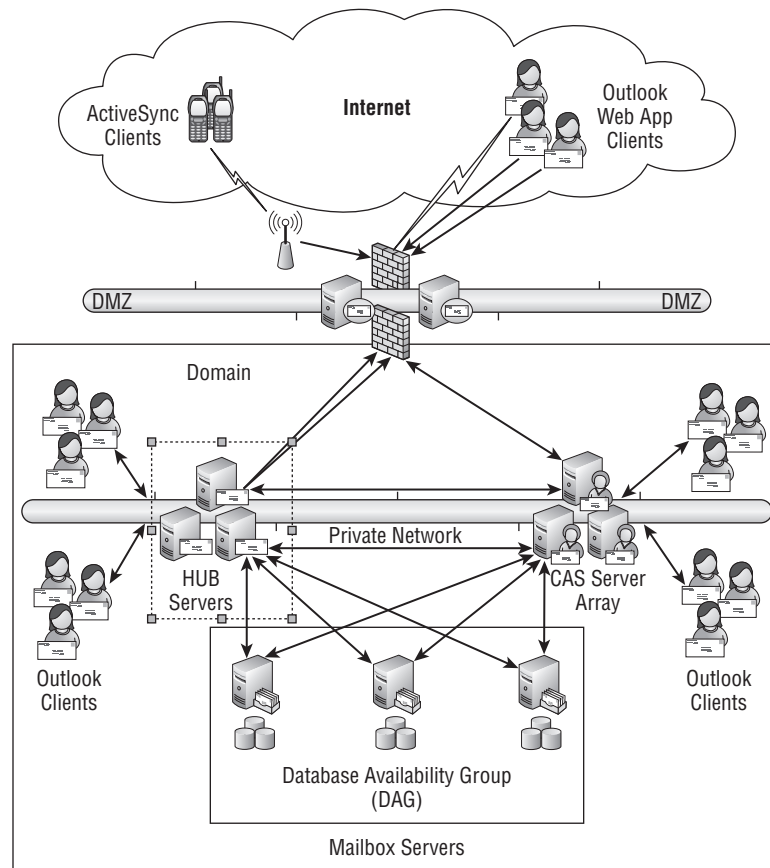
Enhancements to the Hub Transport server role in Exchange Server 2010 include the following:

- ◆ Shadow redundancy provides redundancy for messages the entire time that they are in transit.
- ◆ Write I/O is reduced by about 50 percent.
- ◆ The transport server becomes stateless.
- ◆ The Hub Transport dumpster has been enhanced to use a database replication feedback mechanism to control what messages remain in the dumpster.
- ◆ Messages are deleted from the dumpster once it has been replicated to all database copies.
- ◆ Larger messages are supported without triggering back pressure.

In Exchange Server 2007, this server role is fairly simple to virtualize, but you must take into consideration the disk I/O requirements for the Hub Transport server, because it maintains the queue database. We recommend using RAID 1 or RAID in larger environments. With Exchange Server 2010, this server role becomes easier to virtualize because of the reduction in disk I/O and the stateless database. It is safe to utilize RAID 5 for the database, and we recommend that you deploy at least two Hub Transport servers in every site that hosts an Exchange server.

Figure 6.3 shows a deployment of three Hub Transport servers on the internal or private network. This will facilitate load balancing as well as redundancy for the Hub Transport server.

FIGURE 6.3
Exchange
2007/2010
Hub Transport
server role



Mailbox Server Role

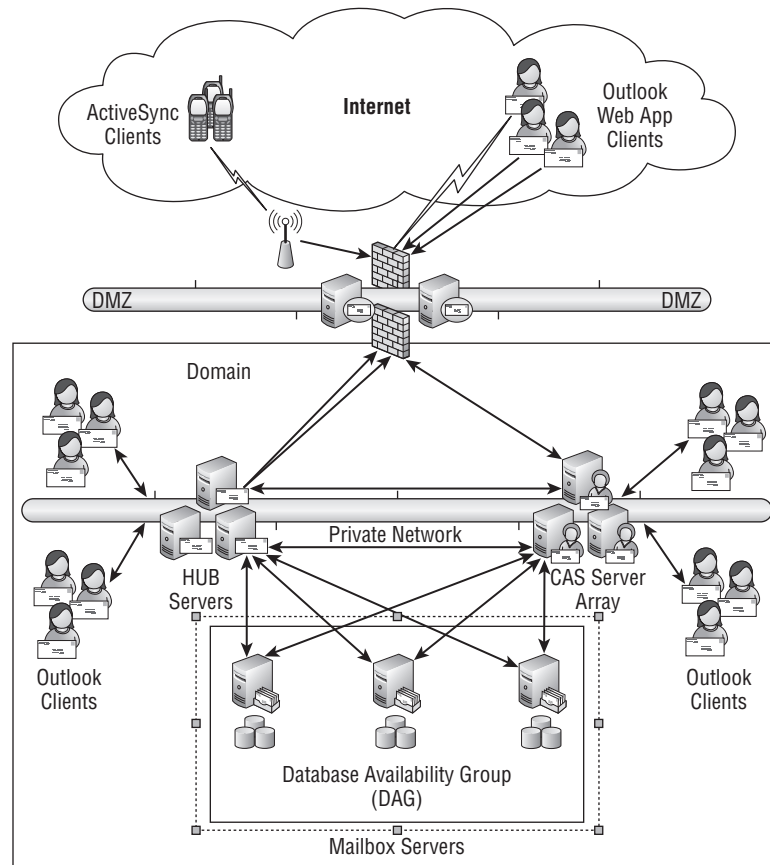
This server role will be the heavy lifter of your Exchange 2010 deployment. Its sole responsibility is to play host to the mailbox and public folder databases. In Exchange 2007 and 2010 Server editions, the public folder is optional and has been provided for backward compatibility.

All connections with the exception of the public folders are accomplished through the Client Access server role. Connections to the public folder databases are still handled by the RPC

endpoint on the Mailbox server role, so clients must still connect directly to the Mailbox server role to access the public folders.

Figure 6.4 shows three Exchange 2010 Mailbox servers deployed as a database availability group (DAG). We will get into DAGs later in the chapter.

FIGURE 6.4
Exchange Server
2010 Mailbox
servers



The Mailbox server role is also the logical unit in a high-availability configuration. Because the Mailbox role performs the database hosting role only, you must have installed a Client Access server and a Hub Transport in any Active Directory site that will host Exchange messaging services.

As mentioned earlier, the CAS role will handle all client connections for mailbox access, and the Hub Transport role will route message throughout the Exchange organization.

New for the Exchange 2010 platform, the concept of the storage group has been deprecated. The term database has been redefined to mean an Exchange database and its transaction log files; they are now thought of as a single unit. In the Exchange 2010 Enterprise edition, each Mailbox server role can host up to 100 databases.

The Exchange 2010 Standard edition can host only five databases per server, but now you can also use DAGs on the Standard edition when installed on Windows Server 2008 Enterprise. It is important to note that Exchange Server 2010 Standard edition can be upgraded to Enterprise edition without reinstalling the product. Because DAGs use Windows clustering services, Windows Server 2008 Enterprise edition is required for implementing DAGs.

Exchange databases are no longer pinned to the Exchange server. Because of this fundamental configuration change, all databases created on the Exchange Mailbox role must have a unique name.



Real World Scenario

REAL-WORLD LESSON LEARNED

When first testing the Exchange 2010 platform with the new LoadGen 2010 tool, a colleague of ours created four individual mailbox databases and then copied the databases and renamed them manually to create a total of twelve databases for testing. When he tried to run the LoadGen 2010 tool, he experienced many errors, one of which was a feedback loop.

We were asked to troubleshoot the configuration and after googling for the specific error, we decided to dismount the last eight databases of the twelve and rerun the LoadGen test tool.

To our surprise, all of the errors disappeared, and we were able to run a successful test. We then deleted the last eight databases, created eight new mailbox databases, and proceeded to run the tool again. All tests completed successfully.

This was a lesson learned for our colleague as well as ourselves. Exchange databases must be truly unique, or you will experience undocumented issues. There was no information anywhere regarding this error message.

Unified Messaging Server Role

This provides the ability to access mail messages via phone or computer. This role is not currently supported under virtualization. Although there are some customers actively virtualizing the Unified Messaging (UM) server role with success, Microsoft explicitly states that it will not support this role virtualized and recommends that you deploy this role on physical servers. You should install the UM server role in an Active Directory site after the Hub Transport role is installed.

Not much testing has been done internally to prove whether this role is a viable candidate for virtualization. Apart from this, we will discuss only some of the capabilities of this role under the Exchange 2010 platform. These are some of the features and enhancements for Exchange 2010:

Auto Attendant Allows for the creation of custom menu prompts to navigate the UM system. The user can then use their voice or touchtone keypad to navigate the system.

Call Waiting light Alerts user of waiting messages.

Subscriber Access Allows a user access to their mailbox, calendar items, and messages using a regular phone line.

Text-to-Speech Allows user to listen to messages in their Outlook mailbox.

To date Microsoft does not support the Unified Messaging server role virtualized because of heavy demands on server CPU and memory. With the introduction of vSphere, however, you can have an eight-vCPU 255 GB memory virtual machine. With these new virtual machine capabilities, some customers have started to virtualize this role. We recommend that you follow the Microsoft Support guidelines about virtualizing Exchange Server 2010 to ensure that you get the best support that Microsoft has offer for virtualized Exchange deployments.

These roles could be deployed on a single server or on multiple servers for modularity. Even though you could install your Exchange 2003 server roles onto multiple servers, you would still get all the Exchange Server services installed and started on each of the server roles. This became a real point of frustration with Exchange and Windows server administrators when trying to secure the various server roles.

Exchange Server 2003 and prior versions were all 32-bit environments, and the maximum addressable memory that Exchange Server 2003 and prior versions could handle was 4 GB. You could increase this limit by using various boot .ini switches, but under heavy usage scenarios, you could exhaust the Windows kernel resources and fragment the usable memory address space. The only way to recover from these issues was to reboot the server. Table 6.1 describes some of the differences among the various Exchange Server versions since Exchange Server 2003.

TABLE 6.1: Comparison of Exchange Server Versions

EXCHANGE SERVER 2003	EXCHANGE SERVER 2007	EXCHANGE SERVER 2010
32-bit environment.	64-bit environment. 32-bit version (not supported in production).	64-bit environment.
4 GB max addressable memory.	Multiterabyte max addressable memory.	Multiterabyte max addressable memory.
About 932 GB addressable database cache.	Multigigabyte database cache for better performance.	Multigigabyte database cache for better performance.
4 storage groups	50 storage groups	Storage groups removed from Exchange 2010. Databases and logs still remain, but both are now referred to as just a database. 100 databases per server max.

TABLE 6.1: Comparison of Exchange Server Versions (CONTINUED)

EXCHANGE SERVER 2003	EXCHANGE SERVER 2007	EXCHANGE SERVER 2010
Storage I/O very heavy because of small block size (about 4 KB) and database cache. Very difficult to scale beyond 2,500 users per Mailbox server and maintain large mailboxes and good end-user experience.	Larger database cache, increased block size and write coalescing to increase performance, and reduction in storage I/O to about 50 to 70 percent.	Schema changes along with sequential database access and large database cache allow for even more storage I/O reduction (around an additional 50 to 70 percent) and larger mailboxes.
Administrative groups used to delegate control and access to Exchange organization. Routing groups used to send and receive mail to nonlocal Active Directory sites.	Routing changes in Exchange 2007 and the introduction of the Hub Transport server role. With Exchange Server 2007, routing is now performed using Active Directory sites.	Routing remains pretty much the same with enhancements to the Hub Transport server role. Shadow messaging is introduced. Shadow messaging ensures that messages are delivered to each hop in the delivery path. If delivery fails anywhere in the delivery path, the message is re-sent from the previous hop to where the message failed. Role-Based Access Control (RBAC) is introduced in Exchange 2010. RBAC provides the ability to granularly grant access permissions to the Exchange 2010 organization.
	Active Directory security groups are used to grant access and delegate access to the Exchange 2007 organization.	The database page size has been increased to 32 KB.
		Database mobility means that databases are no longer pinned to a specific server.
		DAGs allow you to maintain multiple copies of your Exchange databases and provide immediate failover in the case of server or database corruption.
		Microsoft removed the server as the dependency for the Exchange Server databases. Databases have been moved to the organization level of the Exchange Server.

These are just some of differences that make Exchange 2007 and 2010 extremely virtualization friendly. There are whole hosts of other enhancement and improvements, and we invite you to examine these features and improvements to the Exchange 2007 and 2010 platforms.

TIP You can find out more about Exchange Server on the Microsoft Exchange Server home page at www.microsoft.com/exchange/.

Exchange High Availability

Exchange Server high-availability options and features have changed with the evolution of the Exchange product. Microsoft went from stand-alone servers to two-node clustering, and over time and new releases of the Exchange product, it offered four-node clustering.

Exchange Server 2003 and the Windows Server product increased the clustered node numbers to eight nodes, with a seven-active-and-one-passive clustered configuration. Although clustering could provide server and service availability, it was still not considered by the industry as a true highly available solution.

In a true highly available solution, you will find the following characteristics:

- ◆ Server or hardware availability
- ◆ Service availability
- ◆ Data redundancy and availability

Hardware Availability

In clustered solutions, you will encounter a multiple of servers, usually configured identically, all working together to increase the uptime and availability of the application or infrastructure services configured to run as a cluster service or application.

In these scenarios, there is usually a resource DLL that is registered as a cluster resource upon installation of the infrastructure service or application unless the application is a generic application that you are trying to cluster. In this case, you are creating a generic clustered service or application.

In the event of hardware failure or outage, the remaining nodes of the cluster will provide service availability by performing a failover to one of the surviving nodes of the cluster. In most large Exchange Server failover clustering scenarios, it is typical to find 6+2, 5+3, and 4+4 active-passive scenarios.

Service Availability

In the area of service availability, you will want to ensure that in the face of some business disruptive event, the Exchange Server services will be able to provide continued services to the end users without extended downtimes.

Through the use of the clustered resource, the service or services will be constantly monitored for availability, and in the event of some business disruptive event, the service or services will be started on a surviving node of the failover cluster.

In some clustered application scenarios, an additional set of identical services is already running on additional nodes within the cluster. In this scenario, the actual application or infrastructure service recovery time is shorter.

You encounter this scenario in the Exchange Server 2007 clustered continuous replication (CCR) configuration. A subset of the services is already running on the second clustered node, providing the replication services as well as playing replicated transaction logs in the second copy of the Exchange Server database or databases.

Data Redundancy and Availability

In any true highly available infrastructure or application solution, there must be data redundancy. Your clustered solutions are useless if you or your end users cannot access their data (their data being their mailboxes). Data redundancy is one area where traditional failover clustering actually fails.

In the traditional failover clustering scenario, you have multiple clustered nodes accessing a single copy of data stored on shared storage. This is known as a single copy cluster (SCC).

In this scenario, a single node of the cluster is in control of the shared storage at any time. When a node of the cluster fails or becomes unavailable, control of the data on the shared storage is transferred to a surviving node in the clustered scenario.

In the event of logical corruption, the data itself will become unavailable to end users, and recovery procedures will need to be initiated. This could be a simple volume shadow restore that in most cases could take anywhere from minutes to hours depending on the nature of the logical corruption.

This also brings to light the fact that if your clustered solution is locally deployed or deployed using a data center service model, where all nodes of the cluster are local to a single data center, failure of the data center will cause a complete cluster outage rendering services unavailable to all users.

With the introduction of Exchange Server 2007, Microsoft introduced a clustered solution known as clustered continuous replication (affectionately known as CCR). In the CCR scenario, you have a two-node majority node set (MNS) cluster. We will not go into great detail on MNS because it is outside the scope of this book.

With CCR, you have an active and passive cluster node that does not utilize a shared clustered disk model. Quorum is achieved through local disks on each clustered node as well as a file share witness (FSW). The file share witness is a file share usually placed on the Exchange Server 2007/2010 Hub Transport, although this file share could actually be placed on any server.

TIP Microsoft recommends that the FSW be placed on the Exchange Server Hub Transport server and have redundant network connectivity to all clustered nodes.

In the CCR scenario, you have two copies of the Exchange Server databases that are kept in synchronization through asynchronous replication provided at the Exchange Server 2007 application level. Through the use of CCR, you have two copies of the Exchange mailbox data.

This fulfills the data redundancy requirement for a true highly available Exchange 2007 deployment. But wait a minute, what about the data availability requirement for a true highly available Exchange Server deployment?

Exchange Server 2007 CCR allows for deployment in long-distance scenarios known as a *site resilience model*. In this model, we deploy one CCR node in one data center and the other node at a remote data center some distance away. Some of the shortcomings of this solution are network bandwidth and network latencies.

You must ensure that your network can support the replication traffic for the implementation of long-distance CCR. Although Exchange transaction logs have been reduced to 1 MB in size, you could have tens of thousands of them traversing your network over the average workday.

Microsoft also provided standby continuous replication (SCR) as a means for fast disaster recovery. In SCR, you replicated transaction logs to a remote site and delayed the playback of replicated transaction logs into the Exchange Server databases.

Enter database availability groups. In Exchange Server 2010, Microsoft introduced a more mature implementation of CCR and SCR. We like to call it SuperCCR, or CCR on steroids.

This creates a total mobility scenario for the Exchange databases. Exchange Server 2010 databases are now Active Directory objects. In the DAG scenario, copies of the Exchange Server 2010 databases are created on servers within a majority node set cluster now known as a DAG for Exchange Server 2010.

Windows Failover Clustering (WFC) functions have been totally hidden from the Exchange administrator. The Exchange administrator creates the DAG and adds Exchange Server 2010 servers to the DAG as members. In all actuality, you are creating a Windows Server majority node set failover cluster shell, if you will, and then adding nodes to the cluster.

There is now a new Exchange Server 2010 service that manages state of the active and passive copies of the databases. This is the Exchange Server 2010 Active Manager service, which runs on all Exchange Server 2010 servers that are participating in a DAG.

There are two incarnations of the Active Manager, the primary Active Manager or the (PAM) and the standby Active Manager (SAM).

The PAM runs on the Exchange Server 2010 hosting the active copies of the Exchange databases. It constantly monitors Exchange databases and can detect failures. When a failure is detected, the Exchange Server hosting the active databases asks the PAM to perform a failover if the server is still online.

If the server hosting the PAM is unavailable, the PAM role will be transferred to another Exchange server, and the failed databases will be mounted. The PAM also owns the cluster quorum resource and is responsible for making decisions about topologies changes within the DAG. The PAM updates the RPC Client Access service with details about the newly mounted databases so that client connections can be redirected to the server that now hosts the active database.

The SAM role assumes the standby role of monitoring the state of the Exchange database copies and receives requests from the PAM to activate healthy copies of the Exchange databases in the event of a mounted database or server failure.

TIP For more information on the Exchange Server 2010 Active Manager, please see the Exchange Server 2010 online documentation on the TechNet website (<http://technet.microsoft.com/en-us/library/dd776123.aspx>).

Workload Considerations When Virtualizing Exchange Server Roles

Virtualization has made it increasingly simple to deploy systems quickly and with very little effort. Unfortunately, this can also lead to very little thought being put into the performance characteristics of the systems you are deploying. Your virtual data center has a pool of processing power, memory, disk, and network. It's very easy to carve that up to suit the needs of your virtual machine, but what you may not be aware of is the number of resources you are tapping into that are shared by other workloads. This is why proper planning, testing, and tuning are especially important in a virtual environment.

In the physical world, you are able to isolate workloads, at least to a single physical server, if not across the network and storage layers. Although this ensures that you are getting the dedicated resources to the applications being hosted on that single piece of hardware, it leads to a large consumption of power, cooling requirements, and under-utilized hardware. This is largely why the industry is moving toward deploying more and more apps to virtualized hardware. How do we do so in a way that we can guarantee our applications will be as responsive as they were on their over-provisioned hardware?

- ◆ Know the performance characteristics of your application.
- ◆ Gather performance data from the current production environment; if data is not available, use recommended “building blocks” as a starting point.
- ◆ Inventory your current resource consumption to ensure there is proper capacity to move forward and to ensure you are not excessively over-committing your resources.
- ◆ Test using expected workloads in stage or preproduction environment.
- ◆ Provide adequate time to tweak virtual hardware if needed.

With this data in hand, you will be much better prepared to begin planning the production environment, and you will have a sense of how your environment will perform. As Exchange has evolved over the past few years, functions have been moved around, and requirements have changed dramatically. In the early days of Exchange, we required bridgehead servers to send and receive mail. That requirement was somewhat lifted with Exchange 2000 and 2003 as the SMTP service was incorporated into the Mailbox servers. Now with Exchange 2007 and 2010, we again have the concept of a bridgehead server in the form of the Edge and Hub Transport roles. This is just one of many changes that have occurred in the past few years. Needless to say, if you haven't touched Exchange since 2000 or 2003 (or even 5.5), chances are you are unfamiliar with the new performance characteristics of the latest versions. This section will help introduce you to some of the basic fundamentals for appropriately sizing each role individually and for combining some of these roles for smaller environments.

Edge Transport

As you learned in the previous section, the Edge Transport role is designed to be the entry and exit points for all mails entering your organization. This role is deployed on the perimeter network and provides a secure mechanism for providing SMTP relay and smart host services by presenting a minimal attack surface and not participating in your organization's Active

Directory forest. The Edge Transport role is the ideal location for scanning, filtering, and tagging messages based on policies set by your organization. This is done by employing transport agents that act on messages based on content, attachments, or virus and spam signatures.

PROCESSOR

The Edge Transport role is able to take advantage of a multiprocessor configuration of up to 12 processor cores. This is the maximum recommended configuration for the Edge Transport role. The maximum recommended configuration is not a maximum supported configuration; rather, it's a balance between performance and cost. In most environments, four to eight processors will be the recommended configuration to be used by the Edge Transport because most people will choose to scale out their environments to multiple smaller virtual machines rather than scale up. Scaling out allows for greater redundancy and can lead to greater performance in the virtual environment because of being able to spread the load across more vSphere hosts.

When planning processor configurations for virtualized Edge Transport servers, there are some key factors that will contribute to processor utilization:

Message rate How many messages must you be able to process per second?

Message size What is the average size of messages that are processed?

Transport agents What types of transport agents will be installed and enabled? This includes antivirus, spam, content filters, and so on.

The Edge Transport role is more difficult to size than the Hub Transport role because of the level of message hygiene and inspection that happens at this layer. Being the point of entry into most organizations, chances are this role will handle scanning for spam and viruses and will validate message recipients. This means that every message needs to be evaluated and acted upon. Table 6.2 provides a baseline that you can use to compare against your organization.

TABLE 6.2: Edge Transport Performance Metrics: Internal Microsoft Deployment

PERFORMANCE METRIC	VALUE
SMTP Connections/Sec	55
% Connections Accepted	80 percent
SMTP Messages IMF Scanned/Sec	3.7
% SMTP Messages Passed IMF Scanning	80 percent
SMTP Messages A/V Scanned/Sec	3
Avg. Message Size	70 KB
CPU Utilization	20 percent

Microsoft TechNet: <http://technet.microsoft.com/en-us/library/dd346701.aspx#EdgeMemory>

Memory utilization for the Edge Transport role is minimal, relatively speaking. Significant memory use by the Edge Transport role is usually attributed to the database cache for the queue database, scenarios involving large queues, and EdgeSync requirements. The recommended configuration for memory allocation to the Exchange 2007 Edge Transport role is 1 GB of memory per every virtual CPU, with a minimum requirement of 2 GB. Following this recommendation, an Edge Transport virtual machine with one virtual CPU would require at least 2 GB of memory and every additional CPU allocated above has a recommendation of one additional gigabyte or memory per CPU. Exchange 2010 has increased the minimum requirement to 4 GB and maintains the 1 GB per CPU recommendation.

A factor to consider when dealing with large organizations or when planning for large queue scenarios is the amount of memory that will be required to hold each message in memory. Since queue messages and recipient information are held in memory, you need to be sure you are adequately sized to deal with the large queue. Table 6.3 provides the memory overhead used to calculate potential memory usage for Exchange 2007 and can be used as a baseline for Exchange 2010 transport servers.

TABLE 6.3: Memory Overhead Factor

MEMORY FACTOR	MEMORY ALLOCATED
Per message	3KB
Per recipient	1KB

Microsoft TechNet: [http://technet.microsoft.com/en-us/library/bb738124\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb738124(EXCHG.80).aspx)

STORAGE

Although not as storage intensive as the Mailbox role, the Edge Transport role needs to be able to provide a certain level of I/O responsiveness to perform well. When deploying on a virtualized platform, it is easy to take for granted the importance of the underlying storage design. To the vSphere administrator, a VMFS partition with a few hundred gigabytes of free space may seem like the ideal location for your Edge Transport's storage, but let the mail start flowing, and you may realize that the underlying physical storage design is not up to the task. Designing the correct storage environment for your virtual Edge Transport servers is no different from how you would design for a physical environment. You need to understand the I/O and space requirements to sustain the unexpected queue buildup that can occur from time to time.

The amount of space required mostly depends on how large we expect the queue database to grow and how much logging we want to enable and keep. The following should be considered when planning how much space to allocate to an Edge Transport server; this does not include space for the operating system and page file:

Message tracking logs The size will depend on the amount of mail processed by your organization. If you already have an Exchange deployment, you can determine the current log generation rate. How much data is kept is configurable by setting the number of day's worth of logs to keep.

Protocol logs Sizes vary based on mail activity but can be configured with a maximum directory size and maximum age.

Connectivity logs Same as protocol logs.

Transport agent logs Same as protocol logs but may vary depending on the agent.

Queue database The amount of space taken up by the queue database is determined by the average message size multiplied by the maximum queue allowable. As an example, if you are designing the Edge Transport to hold a maximum queue of 100,000 messages at 50 KB per message, you must allocate 5 GB to the queue database.

Transaction logs Space is generally not a concern with the Edge Transport transaction logs because the depth is limited by the use of circular logging.

The Edge Transport role must have at least 500 MB (Exchange 2007 SP1) of free space (4 GB if you are planning on running Exchange Server 2007 RTM) to avoid having Exchange activate the back pressure feature.

Exchange 2010 has added some intelligence into the free disk space requirements for the queue database and transaction logs for transport-related ESE databases. In Exchange 2010, back pressure monitors a percentage based on the disk size since a static number like 5 percent is relative to the disk space available. The formulas are as follows:

- ◆ Free hard disk space for message queue database

$$100 * (\text{hard disk size} - \text{fixed constant}) \div \text{hard disk size}$$

$$\text{Fixed constant} = 500 \text{ MB}$$

The result is represented as a percentage of consumed hard disk space and is the high level. These results are rounded down to the nearest integer. The medium level, by default, is the high level minus 2 percent, and the normal level is the high level minus 4 percent.

- ◆ Free hard disk space for message queue database transaction logs

$$100 * (\text{hard disk size} - \text{MAX}(5 \text{ GB or } 3 * \text{DatabaseCheckPointDepth})) \div \text{hard disk size}$$

The DatabaseCheckPointDepth parameter is defined in the EdgeTransport.exe.config file located in the <Install Dir>\Bin directory.

The result is represented as a percentage of consumed hard disk space and is the high level. These results are rounded down to the nearest integer. The medium level, by default, is the high level minus 2 percent, and the normal level is the high level minus 4 percent.

BACK PRESSURE

Back pressure is a monitoring feature introduced in Exchange 2007 that allows the system to monitor its resources, including disk space, memory, and CPU, and make temporary changes to allow for the excessive resource consumption to pass, usually a large queue. Exchange 2007 reacts to back pressure by rejecting incoming connections. Exchange 2010 accepts all incoming connections, but messages over those connections are accepted at a slow rate or rejected until resource pressure is relieved.

Performance of the storage provided to the Edge Transport role is essential to ensure efficient mail flow in and out of your organization. Understanding the I/O characteristics of the Edge Transport role is key to designing a proper storage solution.

The Edge Transport role will attempt to store incoming mail in memory and in the transaction logs to minimize writing to disk. However, if memory is low, only a portion of each message will be stored in memory and in the transaction logs; the rest of the message must be written to the database. For this reason, it is important to keep the database on properly sized and performing disks, whether you choose virtual disks or raw device disks (pass-through).

Along with the queue database, the temp directory must have well-performing disks and is usually placed on the same volume as the queue database. The temp directory is used during the content conversion process, as messages come in that require conversion, the data is streamed to the temp directory to be processed. The location of the queue database and temp directory is stored in the `EdgeTransport.exe.config` file in the Exchange install directory.

To change the default location of the queue database and temp directory, follow these steps:

1. Stop the transport service by issuing the following command: `net stop MExchangeTransport`.
2. Navigate to the `%programfiles%\Microsoft\Exchange Server\Bin` directory (`%programfiles%\Microsoft\Exchange Server\v14\Bin` in Exchange 2010).
3. Locate and open the `EdgeTransport.exe.config` file in a text editor such as Notepad.
4. The `<appsettings>` section in the `EdgeTransport.exe.config` file contains the configurable settings of the edge transport process. The tags are as follows:

`QueueDatabasePath`: Queue database location

`QueueDatabaseLoggingPath`: Queue database transaction logs

`TemporaryStoragePath`: Temp directory

5. Start the transport service by issuing this command: `net start MExchangeTransport`.

```
<configuration>
<runtime>
<gcServer enabled="true" />
</runtime>
<appSettings>
<add key="QueueDatabasePath" value = "C:\Program Files\Microsoft\Exchange
Server\Queue" />
<add key="TemporaryStoragePath" value = "C:\Temp" />
<add key="QueueDatabaseLoggingPath" value = "C:\Program Files\Microsoft\
Exchange Server\Queue" />
...
</appSettings>
</configuration>
```

When designing the storage system for your Edge Transport servers, you will need a baseline I/O profile to work with. If you are unable to use data from your production environment,

you can start with numbers published on Microsoft TechNet's website ([http://technet.microsoft.com/en-us/library/bb738141\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb738141(EXCHG.80).aspx)). Although originally published for Exchange 2007, they provide a good baseline when the expected load is unknown. These numbers include an average message size and the message rate. For the example in Table 6.4, a message size of 60KB is used with an expected message rate of 20 messages per second.

TABLE 6.4: Disk IOPS per Message (60 KB)

EDGE TRANSPORT DATABASE I/O	APPROXIMATE EDGE TRANSPORT I/O
Log write I/Os per message (sequential)	11
DB write I/Os per message (random)	4.5
DB read I/Os per message (random)	2.5
Total I/Os per message	18

Microsoft TechNet: [http://technet.microsoft.com/en-us/library/bb738141\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb738141(EXCHG.80).aspx)

Using Table 6.4 and the message rate (20 messages per second), you can calculate the database and transactional I/O required to provide the performance necessary. In the following formulas, 20 percent has been added to provide for growth and traffic spikes.

$$\text{DB IOPS per message} = (20 \times (4.5 + 2.5)) + 20\% = 168$$

$$\text{LOG IOPS per message} = (20 \times 11) + 20\% = 264$$

As you can see, the required database IOPS came in at 168, with the log IOPS at 264. With the I/O and sizing requirements in hand, you can work with your storage team to design the best disk layout for your Edge Transport virtual machines.

Client Access Server

The Client Access server role has evolved from what used to be called the front-end server in Exchange 2003. At that time, the front-end server was nothing more than a landing spot for client connections that were then proxied back to the Mailbox server. This meant very little resources were required on these front-end servers since all the heavy lifting was still being done by the Mailbox server, which in most cases had sufficient resources to handle the load.

Exchange 2007 was released with dedicated roles for the various functions required in the Exchange organization. Among them was the CAS role, which as we know now, is a more mature version of the front-end server. The CAS role is intended to be the connection point for all client communications. Exchange 2007 accomplished this, for the most part. Outlook Web Access, Outlook Anywhere, POP, IMAP, Activesync, and Autodiscover connections all terminated with the CAS server from which connections were proxied back to the Mailbox server if required. The main difference now is where things such as content conversion, OWA page rendering, and client configuration were being processed. Those who thought of the CAS as

nothing more than an updated front-end server may have been a bit surprised when they realized that the CAS servers they deployed were under-provisioned and required either extra CPU, memory, or simply additional members to be added to the CAS pool.

The new CAS role was a great leap forward from the old front-end servers, but there was still one client connection point that was missing. MAPI connections from Outlook users within an organization's firewall were still terminating at the Mailbox servers. Exchange 2010 has solved this with RPC client access. Outlook clients connecting to an Exchange 2010 mailbox will now have to connect through the CAS server via RPC client access. This change improves the client experience by reducing the time it takes a client to reconnect after a failover. Instead of taking up to 15 minutes to reconnect, a client will now connect to their mailbox within 30 seconds of a failover.

With this new responsibility, you will need to make sure the CAS role is properly sized to accommodate the increase in load that will be placed by MAPI clients. In the following sections, you will be presented with the data required to size the CAS role appropriately and will be able monitor the role to find any potential bottlenecks.

PROCESSOR

With the CAS server taking on more responsibilities in the latest versions of Exchange Server, you will have to do some planning around processor configuration. In the default installation of Exchange, your CAS servers will be responsible for providing and rendering the Outlook Web Access (or Outlook Web App in Exchange 2010) interface, providing Autodiscover and Availability functions, and being the connection point for Outlook Anywhere (formerly RPC over HTTPS). If you are deploying in an Exchange 2010 environment, the default installation will include RPC Client Access. Finally, you still have the ability to provide Post Office Protocol (POP) and Internet Message Access Protocol (IMAP) capabilities to your end users, but this all requires more processing power.

Microsoft has again done a great job of laying down some initial baselines for Exchange administrators to use when planning their environment. With all the offloading of processing that has been done, the CAS role is able to efficiently use four virtual processors in environments where there is substantial non-MAPI client traffic. You also have the option to deploy in a scaled-out approach, which can provide the required amount of processing capabilities while providing redundancy. The recommended processor configuration for the CAS role is between four to eight virtual processors. The maximum recommended configuration is 12 virtual processors; this is currently more than is supported in a vSphere environment. The consideration on how many virtual processors to assign to a CAS role VM will come by looking at the recommended number of total CAS virtual processors and the number of VMs you desire to scale out to.

For Exchange 2007 and 2010, the recommended mailbox-to-CAS processor ratio is four-to-three. That is, for every four Mailbox role virtual processors, you should deploy three CAS role virtual processors. In an environment with two mailbox VMs, each with four virtual processors, you will require a total of six CAS role virtual processors. How you divide those virtual processors should depend on your scaling-out strategy. You can deploy one CAS role VM with six processors or three CAS role VMs with two virtual processors each. The ratio recommendation assumes all protocols are using Secure Sockets Layer (SSL) encryption and that the processor type and speeds are the same.

MEMORY

As with the increased processing requirement that the new roles have placed on the CAS, the memory usage has increased as well. The CAS role memory usage is mostly based on the number of client connections and the transaction rate at which those clients are processing data. This relationship between memory utilization and client connections is linear; the more client connections and transactions your CAS sees, the more memory you will utilize. As client connections drop, so will the utilized memory. The memory utilization and processor utilization are balanced as well, which means you will see your CAS become processor bound at the same time it becomes memory bound. Table 6.5 lists the current minimum, recommended, and maximum supported memory configurations for CAS servers.

TABLE 6.5: Memory Configurations for Exchange 2007/2010

ROLE/VERSION	MINIMUM SUPPORTED	RECOMMENDED	MAXIMUM PER SERVER
Client Access/ Exchange 2007	2 GB	1 GB per processor (2 GB minimum)	16 GB
Client Access/ Exchange 2010	4 GB	2 GB per core (8 GB minimum)	24 GB

The minimum supported memory configuration must be taken into account when designing your environment. As stated, it is the minimum supported configuration, and it is possible that support will be refused because of a lack of sufficient resources. On the contrary, the maximum per server is a recommendation based on price and performance and does not have any bearing on whether your design is supported.

STORAGE

Because of the mostly stateless nature of the CAS server, disk I/O is not a big concern for this role. As you have read in the previous couple of sections, the CAS role is pretty CPU and memory bound. That's not to say you shouldn't give disk I/O some thought. The basics still hold true; you want to ensure that the OS and page file have the resources they require to avoid a bottleneck. Consider some of the other processes that factor into disk I/O on a Client Access server:

Protocol logging The CAS acts as a web server, a POP server, and an IMAP server, and it provides various other services such as ActiveSync. With all these services, it is recommended that you enable and retain some level of logging to assist in troubleshooting or just to verify that the services are working as expected. This logging creates sequential write activity on the storage and consumes space.

Content conversion If any of your users are IMAP or POP clients, there's a good chance content conversion will come into play in your environment. When messages are sent through the transport engine, they are converted to MAPI and delivered to the mailbox store. When requested by an IMAP or POP client, the mail must be converted to Multipurpose

Internet Mail Extensions (MIME) before being sent back to the client. The CAS is responsible for this conversion. If the message is larger than 64 KB, the conversion takes place on disk, thereby adding I/Os to the storage.

The amount of additional disk I/O placed on the system by logging is minimal, but you should be aware of it. When testing your prototype VMs running the CAS role, be sure to have all the services and logging that you will run in production enabled so that you have a clear indication as to the load presented by your configuration.

Hub Transport

The workload characteristics of the Hub Transport role closely match those of the Edge Transport. Both roles are responsible for sending and receiving mail and applying policies to those messages such as antispam and antivirus filtering. The Hub Transport performs additional functionality that requires consideration when sizing for your environment. The main focus here will be to point out any additional workloads present in the Hub Transport role that are not present in the Edge Transport role. This section, coupled with the Edge Transport section, will provide the knowledge required to deploy transport services that can help meet your organization's requirements.

EDGESYNC

The EdgeSync process uses the pairing of the Edge Transport and Hub Transport services to synchronize directory data. This directory data contains the mail-enabled objects (recipient and distribution lists) as well as configuration data for the Exchange organization. The transfer of data happens at the Hub Transport servers and is pushed out to the Edge Transport servers. To facilitate this transfer, the data is held in memory on the Hub Transport servers. Approximately 4 KB of memory per mail-enabled object is consumed by the EdgeSync process. When sizing the memory for the Hub Transport role, consider the additional memory required based on the number of mail objects that are present in the directory as well as growth expectations.

TRANSPORT DUMPSTER

The transport dumpster was introduced with Exchange 2007 to support the new continuous replication features. This feature has carried over to Exchange 2010 for use with protected databases (those in database availability groups). Because of the asynchronous nature of continuous replication, there is the chance that some log files will be missed in the event of a lossy failover. To mitigate the risk of lost mail during a lossy failover, the transport dumpster retains a copy of each message where at least one of the recipient's mailboxes is stored on a protected database (LCR, CCR, DAG).

The I/O footprint of the Hub Transport with the transport dumpster enabled is significantly increased. The storage that holds the queue database will not only have an increase in write activity, but reads will also occur. In Table 6.6, an average message size of 40 KB is used to show the difference in I/Os per message with and without the transport dumpster enabled. Although originally published for Exchange 2007, these numbers provide a good baseline for an Exchange 2010 design.

TABLE 6.6: Hub Transport I/Os per Message

HUB TRANSPORT SERVER DATABASE I/O (STEADY STATE)	TRANSPORT DUMPSTER ENABLED	TRANSPORT DUMPSTER DISABLED
Total IOPS per message (approximately 40 KB)	17	4
Log write I/Os per message (sequential)	7	2
Database write I/Os per message (random)	7	2
Database read I/Os per message (random)	3	0

Microsoft TechNet: [http://technet.microsoft.com/en-us/library/bb738141\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb738141(EXCHG.80).aspx)

With the numbers from Table 6.6, you can use the same formula that was used for the Edge Transport role to calculate the database and transactional I/O requirements for both scenarios. Plug in the desired message rate, such as 20 messages per second, plus 20 percent for growth and traffic spikes.

Transport Dumpster Enabled

DB IOPS per message (TD enabled) = $(20 \times (7 + 3)) + 20\% = 240$

LOG IOPS per message (TD enabled) = $(20 \times 7) + 20\% = 168$

Transport Dumpster Disabled

DB IOPS per message (TD disabled) = $(20 \times (2 + 0)) + 20\% = 48$

LOG IOPS per message (TD disabled) = $(20 \times 2) + 20\% = 48$

You can see the difference enabling the transport dumpster makes on the overall I/O requirements. In reality, this usually isn't something that sways a decision to use continuous replication. The fact is that the extra load placed by the transport dumpster will usually be easily handled by a redundantly configured storage subsystem with multiple disk spindles. Regardless, it is important to keep this data handy when designing your solution to ensure your Hub Transport can satisfy the requirements.

The transport dumpster queue is part of the Hub Transport queue database (mail.queue). The size of the transport dumpster is an organization-wide configuration parameter. To view the current configuration, use the Exchange Management Shell command `get-transportconfig`.

```
[PS] C:\>get-transportconfig
```

```
ClearCategories           : True
DSNConversionMode        : UseExchangeDSNs
GenerateCopyOfDSNFor     : {5.4.8, 5.4.6, 5.4.4, 5.2.4, 5.2.0, 5.1.4}
InternalSMTPServers      : {}
JournalingReportNdrTo    : <>
```

```

MaxDumpsterSizePerStorageGroup : 18MB
MaxDumpsterTime                 : 7.00:00:00
MaxReceiveSize                  : 10MB
MaxRecipientEnvelopeLimit       : 5000
MaxSendSize                      : 10MB
TLSReceiveDomainSecureList      : {}
TLSSendDomainSecureList         : {}
VerifySecureSubmitEnabled        : False
VoicemailJournalingEnabled       : True
HeaderPromotionModeSetting       : NoCreate
WritingBrandingInDSNEnabled      : True
Xexch50Enabled                   : True

```

By default the Exchange 2007 transport dumpster is enabled and configured for 18 MB per protected storage group with a seven-day retention period. The retention period is the amount of time that messages are kept in the transport dumpster. Seven days is the recommended setting in Exchange 2007. Exchange 2010 introduces a new way of managing the size of the transport dumpster. Instead of keeping messages for a set period of time, messages are tracked via the replication pipeline to determine which messages have been delivered and replicated. Once a message has been confirmed delivered and replicated, the message is truncated from the transport dumpster.

The transport dumpster should be sized to your environment. The recommended setting for the transport dumpster is 1.5 times the maximum allowed message size. In other words, if your maximum message size limit is 10 MB, it is recommended to set the transport dumpster size to 15 MB. In environments where there are no message size restrictions, the recommendation is to size the transport dumpster to 1.5 times the average message size.

Each Hub Transport server in the Active Directory site will maintain the globally configured transport dumpster size per protected storage group. In other words, if your environment has 10 protected storage groups, 2 Hub Transport servers, and a transport dumpster size of 10 MB, each Hub Transport will dedicate 100 MB of queue database space to the transport dumpster (10 storage groups × 10 MB). This space must be accounted for when sizing the storage for the queue database. Additionally, each storage group will have an aggregate amount of transport dumpster space spread across all Hub Transport servers. In the previous example, each storage group would have an aggregate of 20 MB of transport dumpster space (10 MB of transport dumpster per storage group × 2 Hub Transport servers).

To set the transport dumpster to suit your environment, use the `set-transportconfig` Exchange Management Shell command:

```
[PS] C:\>Set-TransportConfig -MaxDumpsterSizePerStorageGroup 15MB -MaxDumpsterTime 14.00:00:00
```

```
[PS] C:\>get-transportconfig
```

```

ClearCategories                  : True
DSNConversionMode                : UseExchangedSNs
GenerateCopyOfDSNFor            : {5.4.8, 5.4.6, 5.4.4, 5.2.4, 5.2.0, 5.1.4}
InternalSMTPServers              : {}

```

```

JournalingReportNdrTo      : <>
MaxDumpsterSizePerStorageGroup : 15MB
MaxDumpsterTime           : 14.00:00:00
MaxReceiveSize            : 10MB
MaxRecipientEnvelopeLimit : 5000
MaxSendSize               : 10MB
TLSReceiveDomainSecureList : {}
TLSSendDomainSecureList   : {}
VerifySecureSubmitEnabled  : False
VoicemailJournalingEnabled : True
HeaderPromotionModeSetting : NoCreate
WritingBrandingInDSNEnabled : True
Xexch50Enabled            : True

```

Client Access and Hub Transport Combined Role

In many organizations, it may not make sense to have separate Hub Transport and Client Access servers. This is especially true in organizations that are going through a server consolidation exercise, have existing underutilized Hub Transport or Client Access servers, or that are virtualizing these roles. Smaller organizations can provide redundancy and consolidate by combining these roles in a virtual environment. Combined roles have different sizing requirements that are explained in the following sections.

PROCESSOR

If you refer to the guidance for stand-alone Client Access or Hub Transport servers, you can see that the processor sizing for these roles is mostly calculated as a ratio to the active mailbox processors in the site. When combining these roles, the ratio becomes a simple one-to-one ratio, that is, one CAS/Hub Transport combined role processor to one Mailbox role processor. This makes sizing combined roles trivial. A small organization may choose to deploy a 4-vCPU Mailbox role VM to support 2,000 average users and deploy two 2-vCPU combined role VMs to satisfy the processor ratio and provide CAS and Hub Transport redundancy.

As with other roles, Microsoft has minimum, recommended, and maximum processor requirements and recommendations. If you choose to deploy a combined role virtual machine, keep in mind that the minimum requirement to receive support is two vCPUs. As with the Mailbox role, the recommended configuration is four vCPUs. The maximum recommended configuration is 12 vCPUs. Deploying on more than 12 vCPUs is not recommended.

MEMORY

The combined roles do not present any significant changes in the required memory. The general guidance is to configure your dual-role VM with at least 4 GB of memory; this is the minimum requirement. The recommended memory configuration for a virtualized dual-role virtual machine is 8 GB. This aligns with the recommended four-vCPU configuration. The recommended maximum for a combined role virtual machine is 2 GB per vCPU allocated to the virtual machine. If deployed at the maximum size allowed in vSphere, a combined-role virtual machine would have eight vCPUs and 16 GB of memory.

If choosing to deploy combined role servers, keep in mind that if your current environment experiences high usage in the CAS role or Hub Transport role, that spike could negatively affect the other roles on your system. A good example of this is the IMAP service on the Client Access role. When fetching a large folder via IMAP, the MSExchangeIMAP4 service can consume 100 percent of the available CPU resources to process the large fetch request. This can impact the Hub Transport role and cause a queue to build while the IMAP request is processed. If working on a new deployment, it is recommended that you deploy separate role virtual machines and attempt to combine only after realizing the baseline load placed on your systems during normal work hours.

Mailbox Server

Client perception of the messaging environment you support will likely come down to the performance of the Mailbox server role. Operations performed by the end users will trigger activity to the back-end Mailbox server. For this reason, the Mailbox server is considerably the most important role to size correctly.

Exchange 2010 was designed with architectural changes that significantly improved performance. Building on the improvements made in Exchange 2007 such as the 64-bit design, Exchange 2010 continues to make great strides. Database I/Os are now larger and more sequential, and the database page size has been increased from 8 KB to 32 KB. The larger page size reduces disk I/O and improves performance by caching the larger pages in memory.

Sizing the Mailbox server can be tricky, but Microsoft provides great guidance for processor, memory, and storage design. The Exchange team has kept the Exchange Server Role Calculator updated to make this process even easier. By knowing a few key metrics from an existing Exchange environment or using some baseline numbers built into the calculator, you can quickly have the requirements in hand to assist in further planning. When deploying in a virtual environment, the same compute resource requirements exist. Using the sizing guidance, you will be able to properly size your virtual Mailbox servers to provide the expected performance for your environment.

PROCESSOR

Exchange 2010 introduces some new concepts around high availability that require you to put a bit more effort into the processor design of your Mailbox server role. DAGs, along with the concept of supporting both active and passive database copies on the same Exchange 2010 Mailbox server, mean you must take the additional processing required by passive database copies into consideration. Not only is the Mailbox server role responsible for the back-end processing of end-user mailbox data, but a Mailbox server supporting DAGs must also perform the validation and replay of log files as well as maintain the content index for its active and passive database copies. For every additional copy of a database, the per-mailbox CPU requirement must be increased by 10 percent on the host serving the active mailbox database.

Table 6.7 gives an estimate of the database cache, IOPS, and CPU requirements based on the user profile for Exchange 2010. In this case, the user profile is based on messages sent and received per day. CPU requirements in Table 6.7 are listed in megacycles per active and passive mailbox. The general recommendation is that a passive mailbox will require approximately 15 percent of the megacycles as the active mailbox.

TABLE 6.7: Estimated Database Cache, IOPS, and CPU Requirements by User Profile

MESSAGES SENT OR RECEIVED PER MAILBOX PER DAY	DATABASE CACHE PER MAILBOX IN MEGABYTES (MB)	SINGLE DATABASE COPY (STAND-ALONE) WITH ESTIMATED IOPS PER MAILBOX	MULTIPLE DATABASE COPIES (MAILBOX RESILIENCY) WITH ESTIMATED IOPS PER MAILBOX	MEGACYCLES FOR ACTIVE MAILBOX OR STAND-ALONE MAILBOX	MEGACYCLES FOR PASSIVE MAILBOX
50	3	0.06	0.05	1	0.15
100	6	0.12	0.1	2	0.3
150	9	0.18	0.15	3	0.45
200	12	0.24	0.2	4	0.6
250	15	0.3	0.25	5	0.75
300	18	0.36	0.3	6	0.9
350	21	0.42	0.35	7	1.05
400	24	0.48	0.4	8	1.2
450	27	0.54	0.45	9	1.35
500	30	0.6	0.5	10	1.5

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

MEGACYCLES PER MAILBOX

Table 6.7 is based on the Intel Xeon x5470 processors in a 2 × 4 core configuration. These processors have a clock speed of 3.33 GHz per processor core and deliver about 3,300 megacycles of performance throughput. For calculating the required megacycles when using any other type of processor, a megacycle adjustment is required. The adjusted value is then entered into the Exchange 2010 Server Role Calculator. To calculate the adjusted value, follow these steps:

1. Browse to www.spec.org, and click Results.
2. On the Published SPEC Benchmark Results page, locate CPU2006, and click the Simple CPU2006 Result Search link.
3. In the Available Configurations field, select SPECint2006 Rates, and select the Simple radio button.
4. In the Simple request section, select Processor from the drop-down of search criteria, enter the processor type (such as x5550), and click Search.

5. In the results, find the SPECint_rate2006 section, and locate the server and processor you will be using, such as Dell R710 with the Intel x5550 2.67GHz processor.
6. Make note of the result (252 or 31.5 per core), the number of cores (8), the number of chips (2), and the number cores per chip (4).

The baseline system (HP DL380 G5 x5470 3.33 GHz × 8 cores) has a SPECint_rate2006 results value of 150, or 18.75 per core.

7. To determine the megacycles of the R710 platform, use the following formula:

$$\text{Megacycles per core} = ((\text{new platform per core value}) * (\text{total megacycles of new platform per core})) / (\text{baseline per core value})$$

$$(31.5 * 2670) / 18.75 = 4485.6 \text{ megacycles per core}$$

8. This number is then plugged into step 5 of the Exchange 2010 Exchange Role Calculator.

By using this formula, you can determine exactly how many processor cores you will require when sizing your Exchange environment. Using the default of 3,300 megacycles will most likely yield a conservative result since most enterprise-level processors will be able to produce a higher megacycle throughput. If you are planning on deploying Exchange 2007 in your environment, the process is slightly simpler.

In Exchange 2007, the concept of Database Availability Groups does not exist, nor does the concept of replication content index data. This significantly reduces the per-mailbox CPU requirements. Additionally, Exchange 2007 still supports the shared-storage model for high availability known as *single-copy clusters*. The rule of thumb for sizing CPU requirements in Exchange 2007 revolved mainly on mailbox profile, mailbox count per server, and whether any third-party utilities were to be used such as Microsoft Forefront.

From a budgeting perspective, the guidance from Microsoft about mailboxes per core for Exchange 2007 has been 1,000 average users per one CPU core. This model allows for up to 4,000 users on a 4× core server. In a virtual environment, these numbers work out quite well given the size of physical server hardware that is available these days. With a maximum recommendation of 12 CPU cores per Mailbox server, companies that have standardized on higher-end physical servers containing 16 or more CPU cores either have to adjust their standards or have to over-commit hardware. Partitioning the larger physical servers into smaller virtual machine units allows these companies to efficiently use the resources they have purchased.

Table 6.8 provides the current minimum supported recommended, and maximum recommended processor configuration for the Exchange Mailbox role for both Exchange 2007 and 2010. The maximum recommended is not a support consideration, but rather a balance between cost and performance.

TABLE 6.8: Processor Configurations for Exchange Mailbox roles

SERVER ROLE	MINIMUM	RECOMMENDED	MAXIMUM RECOMMENDED
2010 Mailbox	2 × processor core	4 × processor cores	12 × processor cores
2007 Mailbox	1 × processor core	4 × processor cores	12 × processor cores

MEMORY

The guidance about providing memory for Exchange 2010 has changed quite significantly from that given for Exchange 2007. Not only has the maximum usable memory cache per mailbox increased (from 7 MB to 30 MB), but the per-database cache has been increased from 20 MB to 100 MB for databases with multiple copies. This increase in database cache directly translates to reduced disk I/O by storing more database changes in memory, thus providing improved coalescence of disk I/O, which equals overall fewer disk I/Os.

To make calculating the amount of usable mailbox cache available while providing the required memory for the OS and applications, refer to Table 6.8 and Table 6.9. As was the case in previous versions of Exchange, the number of databases (previous versions referred to storage groups) determines the maximum amount of memory required. Table 6.9 provides the minimum memory requirements given the number of databases (active and passive).

TABLE 6.9: Minimum Memory Required per Database (Active and Passive)

DATABASE COUNT	EXCHANGE 2010 MINIMUM REQUIRED PHYSICAL MEMORY
1–10	2 GB
11–20	4 GB
21–30	6 GB
31–40	8 GB
41–50	10 GB
51–60	12 GB
61–70	14 GB
71–80	16 GB
81–90	18 GB
91–100	20 GB

Chances are you will be serving more mailboxes than the minimum requirements listed in the table could handle, but keep the numbers in mind because they are the minimum requirements. Table 6.10 provides the recommended database cache based on the user profile (messages sent and received per day) and the estimated IOPS for a stand-alone database and multiple-copy databases. Note the difference in IOPS between single and multiple databases. This is because of the increased log checkpoint depth target (100 MB) allowed for multicopy databases.

TABLE 6.10: Estimated IOPS Based on Mailbox Profile and Mailbox Database Cache

MESSAGES SENT/ RECEIVED PER MAILBOX PER DAY (~75KB AVERAGE MESSAGE SIZE)	DATABASE CACHE PER USER (MB)	SINGLE DATABASE COPY (STAND-ALONE): ESTIMATED IOPS PER MAILBOX	MULTIPLE DATABASE COPIES (MAILBOX RESILIENCY): ESTIMATED IOPS PER MAILBOX
50	3	0.06	0.05
100	6	0.12	0.1
150	9	0.18	0.15
200	12	0.24	0.2
250	15	0.3	0.25
300	18	0.36	0.3
350	21	0.42	0.35
400	24	0.48	0.4
450	27	0.54	0.45
500	30	0.6	0.5

With the previous tables you can determine the minimum memory requirements based on the number of databases you will house on a single Mailbox server and the recommended database cache per mailbox based on user profile. The final piece of this memory puzzle is to provide adequate memory to the Mailbox server to ensure the required mailbox database cache is satisfied as well as the minimum server requirements. Table 6.11 lists the database cache size provided to the Mailbox server role given the amount of physical memory installed. Both mailbox-only and multirole (Mailbox + Hub Transport) servers are covered in this list.

TABLE 6.11: Mailbox Database Cache Sizes

SERVER PHYSICAL MEMORY (RAM)	DATABASE CACHE SIZE (MAILBOX ROLE ONLY)	DATABASE CACHE SIZE MULTIPLE-ROLE (FOR EXAMPLE, MAILBOX + HUB TRANSPORT)
2 GB	512 MB	Not supported
4 GB	1 GB	Not supported
8 GB	3.6 GB	2 GB
16 GB	10.4 GB	8 GB

TABLE 6.11: Mailbox Database Cache Sizes (CONTINUED)

SERVER PHYSICAL MEMORY (RAM)	DATABASE CACHE SIZE (MAILBOX ROLE ONLY)	DATABASE CACHE SIZE MULTIPLE-ROLE (FOR EXAMPLE, MAILBOX + HUB TRANSPORT)
24 GB	17.6 GB	14 GB
32 GB	24.4 GB	20 GB
48 GB	39.2 GB	32 GB
64 GB	53.6 GB	44 GB
96 GB	82.4 GB	68 GB
128 GB	111.2 GB	92 GB

STORAGE

The latest versions of Exchange have reduced disk I/O by as much as 90 percent compared to Exchange 2003 and previous versions. The release of 64-bit only versions of Exchange allowed for a greatly increased mailbox database cache. As described in the previous section, increased mailbox database cache directly affects disk I/O by reducing the need to go to disk for frequently accessed database pages. Database page size was increased from 4 KB in Exchange 2003 and older to 8 KB in 2007 and now to 32 KB in Exchange 2010; again, all of this helps achieve that reduction in I/O, which helps in designing your storage environment.

Regardless of how much the I/O required per mailbox is reduced, there are still latency concerns, which unless addressed will impact the end user's experience. User profile, mailbox size, third-party utilities, and mobile devices all contribute to the I/O footprint presented to your Mailbox servers. Although Microsoft has done a great job at reducing how many disk spindles are required to provide acceptable disk latency, you don't just want to ask the storage administrators for a few hundred gigabytes of space and be on your way. This is especially true in the virtual environment where there is a tendency to mask the underlying physical resources.

Mailbox Role Performance Characteristics

Understanding what contributes to the I/O generated on the Mailbox role is key to designing the proper storage configuration. Not only do you have to take into consideration the load placed by your end users, but now you have the ability to have multiple copies (up to 16) of a single database by using DAGs. The following list provides the factors you must consider when designing the underlying storage:

Mailbox count The number of users directly impacts the I/O per server and capacity requirements.

Mailbox size Setting a mailbox quota will greatly improve the predictability of required storage and future growth based on mailbox count.

Concurrency Knowing the concurrent number of users will help in determining how much active I/O to plan for. If the server hosts users from differing geographic locations, you may be able to leverage this to more efficiently utilize disk performance.

Mailbox profile The mailbox profile is determined by the amount of mail sent and received per user per day and the average message size. The profile type directly impacts transactional I/O.

Outlook mode Outlook users have the option to use online or cached mode to connect to their mailboxes. Cached mode significantly reduces the I/O load placed on the Mailbox server and allows the users to work offline.

Third-party applications Tools such as client-side search tools may index the online copy of a mailbox or public folders, thus increasing load on the server. Utilities that run against the server such as third-party mobile device applications or antivirus software can also lead to increased I/O.

Data replication Whether the replication is storage based or you choose to use a built-in mechanism such as CCR or DAG, there will be an increase in I/O at the server level.

You can collect this information in numerous ways. If you have an existing environment, you can use tools such as the built-in performance monitor in Windows Server 2008 to gather the counters you want to measure and review the results manually. One tool that simplifies this process is the Microsoft Exchange Server Profile Analyzer. The Profile Analyzer tool collects statistical information from your production mailbox stores and can provide information such as mailbox size, rule count, messages per mailbox, and average message size. For more information, refer to this Microsoft TechNet article:

[http://technet.microsoft.com/en-us/library/bb508856\(EXCHG.65\).aspx](http://technet.microsoft.com/en-us/library/bb508856(EXCHG.65).aspx)

For new environments or environments transitioning from a non-Exchange solution or a platform where the numbers don't exactly translate into usable information, the Exchange team has provided a tool that can be used to size not only the storage but also most aspects of your Exchange Mailbox role. The Exchange 2010 Server Role Requirements Calculator (Figure 6.5) has been the de facto sizing tool used by most Exchange implementers for years. It is developed and maintained by the Microsoft Exchange Team, and who would know best how the Mailbox role should be deployed? The calculator has been updated for Exchange 2010 and should be used even if you have your own performance metrics, at least to validate your design against the recommended best practices.

You can find the Exchange requirements calculator at this URL:

<http://go.microsoft.com/fwlink/?LinkId=178613>

When choosing to deploy a new Exchange environment using replicated databases, whether it be CCR or DAG, it is important to consider the impact replication has on storage. The Exchange Server Role Requirements Calculator takes into account the additional I/O and processing overhead that continuous replication adds to the Mailbox server role. With this data in hand, you and your storage vendor can design the underlying storage design that suits your sizing and performance requirements. Fortunately, virtualization has not changed the design aspects of the underlying storage. How we present and attach the storage is affected by virtualization, and we'll touch on the relevant aspects of this in the next section.

FIGURE 6.5
Exchange 2010
Server Role
Requirements
Calculator

Exchange 2010 Mailbox Server Role Requirements Calculator
Version: 6.3

Legal Information: This is provided "AS IS" with no warranties, and confers no rights. Use of this application is subject to the Terms of Use - <http://technet.microsoft.com/en-us/office/exch2010/Exchange2010.aspx>

Instructions: Fill in the blue variables. Choose the appropriate drop-downs for the red variables. The calculator will do the rest.
Important: This tool should only be used for storage modeling purposes. The example configuration provided within this calculator is just that, an example, and as such, each input option needs to be evaluated as to how it will affect your design. Please consult with your storage vendor regarding the appropriate storage design for your environment and follow recommended

Note1: This calculator assumes that the only role installed on the server is the Mailbox Server Role from a CPU sizing perspective.
Note2: The calculated IOPS value has an accuracy of +/- 20% accuracy and does not include third party products that may generate additional database I/O.
Note3: If third party applications/services will be utilized, please refer to the third party manufacturer to determine if the application/service will have any I/O or capacity impacts on the solution.
Note4: This calculator distributes the different tiers of mailboxes across each database (in other words, mailboxes do not have dedicated databases).

Role Requirements Input Factors - Environment Configuration
Step 1: Please enter in the appropriate information for cells that are blue and choose the appropriate drop-downs for cells that are red concerning your messaging environment's configuration. For optimal sizing, choose a multiple of the total number of database copies you have selected for the number of mailbox servers.

Exchange Environment Configuration	Value
Global Catalog Server Architecture	0 (0)
Server Multi-Role Configuration (MEX-CAS-HT)	No
High Availability Deployment	No
Site Resiliency Deployment	No
Site Resiliency Deployment (Mailbox/Database)	No
Number of Mailbox Servers Hosting Active Mailboxes / DAG	6
Number of Database Availability Groups	1

Exchange Data Configuration	Value
Data Overhead Factor	20%
Mailbox Moves / Week Percentage	1%
Dedicated Maintenance / Restore LUN?	No
LUN Free Space Percentage	20%
Log Shipping Network Compression	Enabled
Log Shipping Compression Percentage	50%

Mailbox Database Copy Configuration	Value
Total Number of Mailbox Database Copy Instances (includes Active Copies within DAG)	0
Total Number of Lagged Database Copy Instances within DAG	0
Mailbox Database Copy Lag (minutes)	0

Lagged Database Copy Configuration	Value
Lagged Log Copy (minutes) (minimum)	0

Exchange I/O Configuration	Value
I/O Overhead Factor	20%
Additional I/O Requirement / Server	0.0%

Ready | Input | Role Requirements | LUN Requirements | Backup Requirements | Log Replication Requirements | Storage | 60%

VMFS vs. RDM

The topic of VMFS partitions and raw device mappings should not be new. Chapter 3 discussed the two briefly in respect to deploying Microsoft Cluster Services and failover clustering in a virtual environment. You'll recall when deploying a shared storage cluster in a virtual environment, the recommendation for enterprise environments was to deploy a cluster using virtual machines across ESX hosts. This required the use of physical mode raw device mappings (RDMs). Clusters aren't the only reason you need to consider deploying raw device mappings.

A common misconception is that RDMs, also described as pass-through storage devices (because the path from the virtual machine to the storage passes through the hypervisor with no virtualization interaction), perform better than a virtual disk file residing on a VMFS. Although this may have been the case in earlier versions of many hypervisors, the VMFS file system has been tuned to the point where latency times between VMFS and RDMs are within microseconds of each other. These days, performance rarely plays a role when deciding whether to deploy on VMFS or RDM.

Deciding on the storage access method to use should be based on your design requirements. Common factors include support, such as in the case of clustering. To deploy a Microsoft cluster, you must use physical-mode RDMs as the shared disks. How you plan to back up the data, in this case the Exchange databases and logs, can affect the method of storage access. Many Exchange deployments rely on the VSS snapshot technology introduced with Windows 2003. Hardware-based VSS solutions from the leading storage vendors often require that a backup agent be installed on the Exchange Mailbox server to initiate the VSS activities. These agents require direct access to the back-end storage devices to achieve the cloning or snapshotting at the storage layer once Exchange has been quiesced. This type of access will, in most cases, rely on physical-mode RDMs.


Raw device mappings come at a price. Because the virtualization layer is removed (in physical-mode RDMS), the ability to clone, migrate, and initiate a VMware Data Recovery backup and then snapshot the entire VM is lost. Migration via VMotions will continue to work, but many of the other features available by having the entire virtual machine being comprised of a set of files are lost.

Although the Exchange mailbox database has been optimized to reduce the load on the underlying storage, there is still plenty to consider when designing a new environment. Fortunately for you, virtualization does not add much to the already complex process. Understanding the storage technical requirements up front, knowing the support considerations, and working with your storage vendors can significantly ease the burden of this cornerstone in your virtualized Exchange design.

Unified Messaging

The Unified Messaging role (UM) was first introduced with Exchange 2007 and has been carried into Exchange 2010. UM allows for the integration of your company's telephony (voice and fax) and email system. UM users are able to seamlessly access their voicemail, access their email, and send and receive faxes using a single client. Unified Messaging servers should be placed in close proximity to the PBX system that they will integrate with and optimally on the same VLAN or broadcast domain as the IP gateway, which will perform the translation to Voice over IP (VoIP). Unless quality of service has been configured on the local-area network to ensure the quality of voice calls traversing the data network, it is possible your organization may employ multiple networks dedicated to either voice or data. In these scenarios, virtualizing the UM role of Exchange may require dedicated networking equipment (both physical and virtual) to avoid having to go through multiple network hops to reach the VoIP network. In a virtual environment, this may look something like Figure 6.6. The virtual machine acting as the UM role has two virtual network adapters, one accessing the data network and the other accessing the VoIP network.

FIGURE 6.6
Dual-homed UM
virtual machine
accessing both
the data and VoIP
networks



Unfortunately, Microsoft has chosen not to support the Unified Messaging role in a virtual environment. Because of this, there has not been much adoption to the concept of virtualizing this role. In the following sections, we have listed the resource recommendations that will directly translate to a virtual environment; however, because of the lack of support, we recommend not virtualizing this role until Microsoft sees it as a supportable configuration.

PROCESSOR

Table 6.12 lists the minimum and maximum processor recommendations for the Unified Messaging role.

TABLE 6.12: Processor Recommendations for Unified Messaging

SERVER ROLE	MINIMUM SUPPORTED	MAXIMUM RECOMMENDED
Exchange 2010 Unified Messaging	2× processor cores	12× processor cores
Exchange 2007 Unified Messaging	1× processor core	6× processor cores

MEMORY

Table 6.13 lists the minimum and maximum memory recommendations for the Unified Messaging role.

TABLE 6.13: Memory Recommendations for Unified Messaging

SERVER ROLE	MINIMUM SUPPORTED	RECOMMENDED
Exchange 2010 Unified Messaging	4 GB	2 GB per core
Exchange 2007 Unified Messaging	2 GB	1 GB per core

STORAGE

The Unified Messaging role, much like the Client Access server, is mostly stateless and requires a fraction of the disk I/O required by the Hub Transport or Mailbox roles. Because voicemails and faxes are kept in the end users' mailboxes, the capacity requirements for the UM servers are minimized. The source of disk I/O for the UM role will mostly come from content conversion and protocol logging that has been enabled. Incoming voicemails are converted from the native recorded format to a compressed format defined by the administrator such as MP3. This conversion requires more processing resources than any other, but having well-responding storage will ensure the conversions are as efficient as possible.

Although not being formally supported in a virtualized deployment, virtualizing the UM role in a test environment will provide for a cost-effective development, learning, and testing solution. Testing configuration updates and OS and application updates in a test bubble before deploying to production is usually a requirement for most organizations' change control process. Having an easy-to-replicate and rollback test environment has proved invaluable for many organizations to achieve this process requirement.

MONITORING PERFORMANCE

During deployment and testing, you will want to verify that the system is running as efficiently as possible. Using tools like Service Center Operations Manager can help aid in monitoring, but for real-time analysis or more precise figure capturing, you will want to take advantage of the built-in performance counters. In Tables 6.14–6.17, you will find the counters you can focus on to determine whether your Exchange server is running as efficiently as possible, whether there is a bottleneck in the system, and, if so, which component may be the culprit (according to Microsoft TechNet article [http://technet.microsoft.com/en-us/library/bb201720\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb201720(EXCHG.80).aspx)). These include processor, memory, storage, and network counters.

TABLE 6.14: Processor Counters

COUNTER	EXPECTED VALUES
Processor(_Total)\% Processor Time Shows the percentage of time that the processor is executing application or operating system processes.	Should be less than 75 percent on average

TABLE 6.15: Memory Counters

COUNTER	EXPECTED VALUES
Memory\Available Mbytes Shows the amount of physical memory, in megabytes (MB), immediately available for allocation to a process or for system use.	Should remain above 100 MB at all times
Memory\Pages/Sec Shows the rate at which pages are read from or written to disk to resolve hard page faults.	Should be below 1,000 on average

TABLE 6.16: Storage Counters

COUNTER	EXPECTED VALUES
Logical/PhysicalDisk(*)\Avg. Disk sec/Read	Should be less than 20 milliseconds (ms) on average.
Shows the average time, in seconds, it takes to read data from the hard disk.	Spikes (maximum values) should not be higher than 50 ms.
Logical/PhysicalDisk(*)\Avg. Disk sec/Write	Should be less than 20 milliseconds (ms) on average.
Shows the average time, in seconds, of a write of data to the disk.	Spikes (maximum values) should not be higher than 50 ms.
Logical/Physical Disk(*)\Avg. Disk sec/Transfer	For healthy hard disks, this counter shows approximately 20 ms. Counter values larger than 20 ms, or with large spikes, indicate a possible hard disk issue (for example, failure or slow speed).
Indicates how fast data is being moved (in seconds). Measures the average time of each data transfer, regardless of the number of bytes read or written.	

TABLE 6.17: Network Counters

COUNTER	EXPECTED VALUES
Network Interface(*)\Bytes Total/sec	For a 100-Mbps network adapter, should be below 6–7 Mbps.
Indicates the rate at which the network adapter is processing data bytes.	For a 1000-Mbps network adapter, should be below 60–70 Mbps.
This counter includes all application and file data, in addition to protocol information such as packet headers.	
Network Interface(*)\Packets Outbound Errors	Should be 0 at all times.
Indicates the number of outbound packets that could not be transmitted because of errors.	

The general counters in the previous tables will provide a baseline of understanding how your Exchange servers are performing at the basic compute level. The list of counters for each role can be quite long, and Microsoft has done a great job of publishing these on its TechNet website at the following URL: <http://technet.microsoft.com/en-us/library/dd335215.aspx>.

Gathering all the data provided by these counters can be a daunting task, not to mention actually taking that data and trying to analyze the results. Luckily, there is a great (and free) tool that you can use to do the heavy lifting and provide some great insight into the performance of your Exchange roles.

Performance Analysis of Logs Tool

The Performance Analysis of Logs (PAL) tool uses VBScript and the Microsoft Log Parser tool (another free download) to input previously collected performance logs and analyze the results against known thresholds. The tool will calculate the results and display them in a simple-to-digest HTML report. The PAL tool provides analysis for most Microsoft tier-1 applications.

To use PAL to analyze your environment, you will require the following components:

- ◆ Log Parser 2.2: www.microsoft.com/downloads/details.aspx?FamilyID=890cd06b-abf8-4c25-91b2-f8d975cf8c07&DisplayLang=en
- ◆ PAL Tool: www.codeplex.com/PAL
- ◆ PerfWiz Replacement XML files: <http://blogs.technet.com/mikelag/archive/2008/05/02/perfwiz-replacement-for-exchange-2007.aspx>

The high-level overview of the process to create a PAL report is as follows:

1. Download and install Log Parser and the PAL tool.
2. Install the PerfWiz replacements, either the role based or full counters.
3. Run the performance counters during peak times or while troubleshooting a performance issue.
4. Launch the PAL wizard.
5. Create the PAL job and execute.
6. Review the HTML file created by the PAL tool.

The PAL tool is a valuable resource to add to your toolkit for quick analysis of your Exchange environment. Using the PAL tool on a regular basis to provide historical baselines along with proactive monitoring with tools such as Microsoft System Center Operations Manager will prove useful when troubleshooting any performance problems in your environment.

Backup and Restore

When looking to protect Exchange data, administrators will focus on the Mailbox role. After all, this is where the majority of your data will reside. Depending on the size of the organization, this may involve backing up tens of terabytes (TB) worth of Exchange databases. When dealing with this much data, it is important to implement a solution that will provide a consistent backup and restore method. This will include being able to complete backup operations without interfering with normal user operations.

As mentioned in prior chapters, virtualizing your environment will provide additional methods for protecting your data, but that doesn't mean that traditional backup methods cannot be used. In fact, and especially for Exchange, the traditional agent-based backup methods continue to be the number-one method for backing up Exchange databases. The reason for this is that as