CHAPTER 3

# Understanding Core Exchange Server 2007 Design Plans

IN THIS CHAPTER

- Planning for Exchange Server 2007
- Understanding Active Directory Design Concepts for Exchange Server 2007
- Determining Exchange Server 2007 Placement
- Configuring Exchange Server 2007 for Maximum Performance and Reliability
- Securing and Maintaining an Exchange Server 2007 Implementation

The fundamental capabilities of Microsoft Exchange Server 2007 are impressive. Improvements to security, reliability, and scalability enhance an already road-tested and stable Exchange platform. Along with these impressive credentials comes an equally impressive design task. Proper design of an Exchange Server 2007 platform will do more than practically anything to reduce headaches and support calls in the future. Many complexities of Exchange might seem daunting, but with a proper understanding of the fundamental components and improvements, the task of designing the Exchange Server 2007 environment becomes manageable.

This chapter focuses specifically on the Exchange Server 2007 components required for design. Key decision-making factors influencing design are presented and tied into overall strategy. All critical pieces of information required to design Exchange Server 2007 implementations are outlined and explained. Enterprise Exchange design and planning concepts are expanded in Chapter 4, "Architecting an Enterprise-Level Exchange Environment."

## Planning for Exchange Server 2007

Designing Exchange Server used to be a fairly simple task. When an organization needed email and the decision was made to go with Exchange Server, the only real decision to make was how many Exchange servers were needed. Primarily, organizations really needed only email and eschewed any "bells and whistles."

Exchange Server 2007, on the other hand, takes messaging to a whole new level. No longer do organizations require only an email system, but other messaging and unified communications functionality as well. After the productivity capabilities of an enterprise email platform have been demonstrated, the need for more productivity improvements arises. Consequently, it is wise to understand the integral design components of Exchange before beginning a design project.

## Outlining Significant Changes in Exchange Server 2007

Exchange Server 2007 is the evolution of a product that has consistently been improving over the years from its roots. Since the Exchange 5.x days, Microsoft has released dramatic improvements with Exchange 2000 Server and later Exchange Server 2003. The latest version takes the functionality and reliability of Exchange to the next level, introducing several major enhancements and improvements.

The major areas of improvement in Exchange Server 2007 have focused on several key areas. The first is in the realm of user access and connectivity. The needs of many organizations have changed and they are no longer content with slow remote access to email and limited functionality when on the road. Consequently, many of the improvements in Exchange focus on various approaches to email access and connectivity. The improvements in this group focus on the following areas:

- **"Access anywhere" improvements**—Microsoft has focused a great deal of Exchange Server 2007 development time on new access methods for Exchange, including an enhanced Outlook Web Access (OWA) that works with a variety of Microsoft and third-party browsers, Microsoft ActiveSync improvements, new Outlook Voice Access (OVA), unified messaging support, and Outlook Anywhere (formerly known as RPC over HTTP). Having these multiple access methods greatly increases the design flexibility of Exchange, as end users can access email via multiple methods.
- **Protection and compliance enhancements**—Exchange Server 2007 now includes a variety of antispam, antivirus, and compliance mechanisms to protect the integrity of messaging data.
- **Admin tools improvements and Exchange Management Shell scripting**—The administrative environment in Exchange 2007 has been completely revamped and improved, and the scripting capabilities have been overhauled. It is now possible to script any administrative command from a command-line script. Indeed, the graphical user interface (GUI) itself sits on top of the scripting engine and simply fires scripts based on the task that an administrator chooses in the GUI. This allows for an unprecedented level of control.
- **Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR)**—One of the most anticipated improvements to Exchange Server has been the inclusion of Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR). These technologies allow for log shipping functionality for Exchange databases, allowing a replica copy of an Exchange database to be constantly built from new logs generated from the server. This gives administrators

the ability to replicate in real time the data from a server to another server in a remote site or locally on the same server.

It is important to incorporate the concepts of these improvements into any Exchange design project because their principles often drive the design process.

## Reviewing Exchange and Operating System Requirements

Exchange Server 2007 has some specific requirements, both hardware and software, that must be taken into account when designing. These requirements fall into several categories:

- Hardware
- Operating system
- Active Directory
- Exchange version

Each requirement must be addressed before Exchange Server 2007 can be deployed.

### Reviewing Hardware Requirements

It is important to design Exchange hardware to scale out to the user load, which is expected for up to 3 years from the date of implementation. This helps retain the value of the investment put into Exchange. Specific hardware configuration advice is offered in later sections of this chapter.

### Reviewing Operating System (OS) Requirements

Exchange Server 2007 is optimized for installation on Windows Server 2003. The increases in security and the fundamental changes to Internet Information Services (IIS) in Windows Server 2003 provide the basis for many of the improvements in Exchange Server 2007. The specific compatibility matrix, which indicates compatibility between Exchange versions and operating systems, is illustrated in Table 3.1.

TABLE 3.1 Exchange Version Compatibility

| Version | Windows NT 4.0 | Windows 2000 | Windows 2003 |
|---|---|---|---|
| Exchange 5.5 | Yes | Yes | No |
| Exchange 2000 | No | Yes | No |
| Exchange 2003 | No | Yes | Yes |
| Exchange 2007 | No | No | Yes* |

** 64-bit SP1 or R2 editions only supported*

### Understanding Active Directory (AD) Requirements

Exchange originally maintained its own directory. With the advent of Exchange 2000, however, the directory for Exchange was moved to the Microsoft Active Directory, the

enterprise directory system for Windows. This gave greater flexibility and consolidated directories, but at the same time increased the complexity and dependencies for Exchange. Exchange Server 2007 uses the same model, with either Windows 2000 Server or Windows Server 2003 AD as its directory component.

Exchange 2007, while requiring an AD forest in all deployment scenarios, has certain flexibility when it comes to the type of AD it uses. It is possible to deploy Exchange in the following scenarios:

- **Single forest**—The simplest and most traditional design for Exchange is one where Exchange is installed within the same forest used for user accounts. This design also has the least amount of complexity and synchronization concerns to worry about.
- **Resource forest**—The Resource forest model in Exchange Server 2007 involves the deployment of a dedicated forest exclusively used for Exchange itself, and the only user accounts within it are those that serve as a placeholder for a mailbox. These user accounts are not logged onto by the end users, but rather the end users are given access to them across cross-forest trusts from their particular user forest to the Exchange forest. More information on this deployment model can be found in Chapter 4.
- **Multiple forests**—Different multiple forest models for Exchange are presently available, but they do require a greater degree of administration and synchronization. In these models, different Exchange organizations *live* in different forests across an organization. These different Exchange organizations are periodically synchronized to maintain a common Global Address List (GAL). More information on this deployment model can also be found in Chapter 4.

It is important to determine which design model will be chosen before proceeding with an Exchange deployment because it is complex and expensive to change the AD structure of Exchange after it has been deployed.

#### Outlining Exchange Version Requirements

As with previous versions of Exchange, there are separate Enterprise and Standard versions of the Exchange Server 2007 product. The Standard Edition supports all Exchange Server 2007 functionality with the exception of the following key components:

- **Greater than 75-GB mailbox store**—Exchange Server 2007 can support only a single database of up to 75GB in size. Pre-Exchange 2003-SP2 Standard Exchange only supported up to 16-GB databases. Organizations with small numbers of users or strict storage limits can use this version of Exchange without problems.

**NOTE**

There is no direct upgrade path from the Exchange Standard Edition to the Enterprise Edition. Only a mailbox migration procedure that can transfer mailboxes from a Standard Edition server to an Enterprise Edition server can accomplish an upgrade. Consequently, it is important to make an accurate determination of whether the Enterprise Edition of the software is needed.

- **Multiple mailbox database stores**—One of the key features of Exchange Server 2007 is the capability of the server to support multiple databases and storage groups with the Enterprise Edition of the software. Up to 50 storage groups and/or 50 databases per server are supported. This capability is not supported with the Standard Edition of the product.
- **Clustering support**—Exchange Server 2007 clustering, including traditional Single Copy Clustering (shared storage) and the new Cluster Continuous Replication (CCR), is available only when using the Enterprise Edition of the software. Support for up to an eight-way active-passive cluster on Windows Server 2003 is available. Microsoft requires at least one passive node per cluster.



## Scaling Exchange Server 2007

The days of the Exchange server "rabbit farm" are gone where it is no longer necessary to set up multiple Exchange server sites across an organization and watch them grow as usage of mail increases in the organization. Exchange 2000 originally provided the basis for servers that could easily scale out to thousands of users in a single site, if necessary. Exchange Server 2003 further improved the situation by introducing Messaging Application Programming Interface (MAPI) compression and RPC over HTTP. Exchange Server 2007 further improves the situation by improving RPC over HTTP (now called Outlook Anywhere) and allowing Mailbox servers to scale upward through 64-bit OS support.

Site consolidation concepts enable organizations that might have previously deployed Exchange servers in remote locations to have those clients access their mailboxes across wide area network (WAN) links or dial-up connections by using the enhanced Outlook 2003/2007 or OWA clients. This solves the problem that previously existed of having to deploy Exchange servers and global catalog (GC) servers in remote locations, with only a handful of users, and greatly reduces the infrastructure costs of setting up Exchange.

## Having Exchange Server 2007 Coexist with an Existing Network Infrastructure

Exchange is built upon a standards-based model, which incorporates many industrywide compatible protocols and services. Internet standards—such as DNS, IMAP, SMTP, LDAP, and POP3—are built in to the product to provide coexistence with existing network infrastructure.

In a design scenario, it is necessary to identify any systems that require access to email data or services. For example, it might be necessary to enable a third-party monitoring application to relay mail off the Simple Mail Transfer Protocol (SMTP) engine of Exchange so that alerts can be sent. Identifying these needs during the design portion of a project is subsequently important.

### Identifying Third-Party Product Functionality

Microsoft built specific hooks into Exchange Server 2007 to enable third-party applications to improve upon the built-in functionality provided by the system. For example, built-in support for antivirus scanning, backups, and unified messaging exist right out of the box, although functionality is limited without the addition of third-party software. The most common additions to Exchange implementation are the following:

- Antivirus
- Backup
- Phone/PBX integration
- Fax software

## Understanding AD Design Concepts for Exchange Server 2007

After all objectives, dependencies, and requirements have been mapped out, the process of designing the Exchange Server 2007 environment can begin. Decisions should be made in the following key areas:

- AD design
- Exchange server placement
- Global catalog placement
- Client access methods

### Understanding the AD Forest

Because Exchange Server 2007 relies on the Windows Server 2003 AD for its directory, it is therefore important to include AD in the design plans. In many situations, an AD implementation, whether based on Windows 2000 Server or Windows Server 2003, already exists in the organization. In these cases, it is necessary only to plan for the inclusion of Exchange Server into the forest.

**NOTE**

Exchange Server 2007 has several key requirements for AD. First, all domains must be in Windows 2000 or 2003 functional levels (no NT domain controllers). Second, it requires that the schema in an AD forest be extended for Windows Server 2003 RTM or R2 editions, and that the schema master domain controller be running either Windows Server 2003 SP1 or R2 edition. In addition, at least one global catalog server in each site where Exchange will be installed must be running Windows Server 2003 SP1 or R2.

If an AD structure is not already in place, a new AD forest must be established. Designing the AD forest infrastructure can be complex, and can require nearly as much thought into design as the actual Exchange Server configuration itself. Therefore, it is important to fully understand the concepts behind AD before beginning an Exchange 2007 design.

In short, a single "instance" of AD consists of a single AD forest. A forest is composed of AD trees, which are contiguous domain namespaces in the forest. Each tree is composed of one or more domains, as illustrated in Figure 3.1.

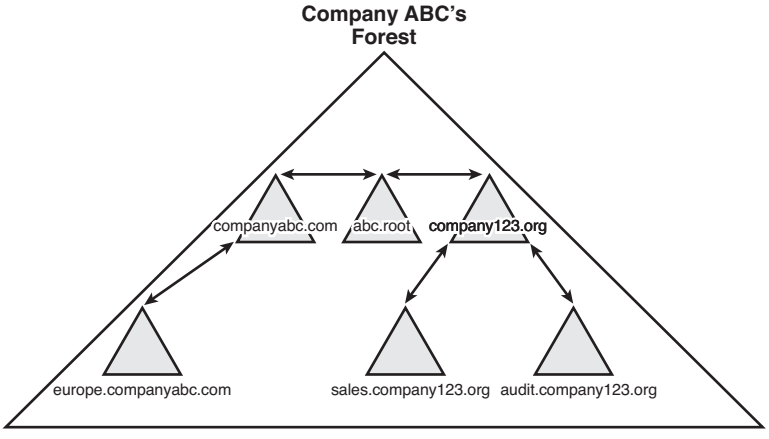


FIGURE 3.1 Multitree forest design.

Certain cases exist for using more than one AD forest in an organization:

- **Political limitations**—Some organizations have specific political reasons that force the creation of multiple AD forests. For example, if a merged corporate entity requires separate divisions to maintain completely separate information technology (IT) infrastructures, more than one forest is necessary.
- **Security concerns**—Although the AD domain serves as a de facto security boundary, the "ultimate" security boundary is effectively the forest. In other words, it is possible for user accounts in a domain in a forest to hack into domains within the same forest. Although these types of vulnerabilities are not common and are difficult to do, highly security-conscious organizations should implement separate AD forests.
- **Application functionality**—A single AD forest shares a common directory schema, which is the underlying structure of the directory and must be unique across the entire forest. In some cases, separate branches of an organization require that certain applications, which need extensions to the schema, be installed. This might not be possible or might conflict with the schema requirements of other branches. These cases might require the creation of a separate forest.

- **Exchange-specific functionality (resource forest)**—In certain circumstances, it might be necessary to install Exchange Server 2007 into a separate forest, to enable Exchange to reside in a separate schema and forest instance. An example of this type of setup is an organization with two existing AD forests that creates a third forest specifically for Exchange and uses cross-forest trusts to assign mailbox permissions.

The simplest designs often work the best. The same principle applies to AD design. The designer should start with the assumption that a simple forest and domain structure will work for the environment. However, when factors such as those previously described create constraints, multiple forests can be established to satisfy the requirements of the constraints.

## Understanding the AD Domain Structure

After the AD forest structure has been chosen, the domain structure can be laid out. As with the forest structure, it is often wise to consider a single domain model for the Exchange 2007 directory. In fact, if deploying Exchange is the only consideration, this is often the best choice.

There is one major exception to the single domain model: the placeholder domain model. The placeholder domain model has an isolated domain serving as the root domain in the forest. The user domain, which contains all production user accounts, would be located in a separate domain in the forest, as illustrated in Figure 3.2.

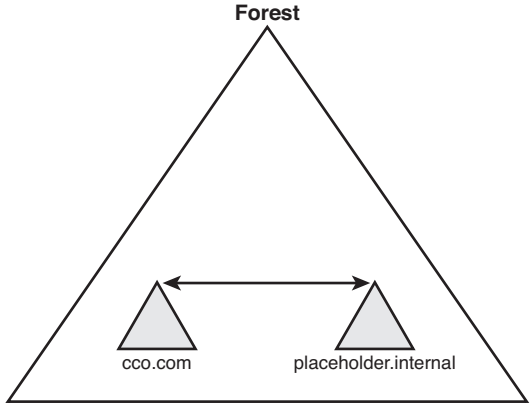


FIGURE 3.2 The placeholder domain model.

The placeholder domain structure increases security in the forest by segregating high-level schema-access accounts into a completely separate domain from the regular user domain. Access to the placeholder domain can be audited and restricted to maintain tighter control on the critical schema. The downside to this model, however, is the fact that the additional domain requires a separate set of domain controllers, which increases the

infrastructure costs of the environment. In general, this makes this domain model less desirable for smaller organizations because the trade-off between increased cost and less security is too great. Larger organizations can consider the increased security provided by this model, however.

## Reviewing AD Infrastructure Components

Several key components of AD must be installed within an organization to ensure proper Exchange Server 2007 and AD functionality. In smaller environments, many of these components can be installed on a single machine, but all need to be located within an environment to ensure server functionality.



### Outlining the Domain Name System (DNS) Impact on Exchange Server 2007 Design

In addition to being tightly integrated with AD, Exchange Server 2007 is joined with the Domain Name System (DNS). DNS serves as the lookup agent for Exchange Server 2007, AD, and most new Microsoft applications and services. DNS translates common names into computer-recognizable IP addresses. For example, the name `www.cco.com` translates into the IP address of `12.155.166.151`. AD and Exchange Server 2007 require that at least one DNS server be made available so that name resolution properly occurs.

Given the dependency that both Exchange Server 2007 and AD have on DNS, it is an extremely important design element. For an in-depth look at DNS and its role in Exchange Server 2007, see Chapter 6, "Understanding Network Services and AD Domain Controller Placement for Exchange Server 2007."

### Reviewing DNS Namespace Considerations for Exchange

Given Exchange Server 2007's dependency on DNS, a common DNS namespace must be chosen for the AD structure to reside in. In multiple tree domain models, this could be composed of several DNS trees, but in small organization environments, this normally means choosing a single DNS namespace for the AD domain.

There is a great deal of confusion between the DNS namespace in which AD resides and the email DNS namespace in which mail is delivered. Although they are often the same, in many cases there are differences between the two namespaces. For example, CompanyABC's AD structure is composed of a single domain named `abc.internal`, and the email domain to which mail is delivered is `companyabc.com`. The separate namespace, in this case, was created to reduce the security vulnerability of maintaining the same DNS namespace both internally and externally (published to the Internet).

For simplicity, CompanyABC could have chosen `companyabc.com` as its AD namespace. This choice increases the simplicity of the environment by making the AD logon user principal name (UPN) and the email address the same. For example, the user Pete Handley is `pete@companyabc.com` for logon, and `pete@companyabc.com` for email. This option is the choice for many organizations because the need for user simplicity often trumps the higher security.

#### Optimally Locating Global Catalog Servers

Because all Exchange directory lookups use AD, it is vital that the essential AD global catalog information is made available to each Exchange server in the organization. For many small offices with a single site, this simply means that it is important to have a full global catalog server available in the main site.

The global catalog is an index of the AD database that contains a partial copy of its contents. All objects within the AD tree are referenced within the global catalog, which enables users to search for objects located in other domains. Every attribute of each object is not replicated to the global catalogs, only those attributes that are commonly used in search operations, such as first name and last name. Exchange Server 2007 uses the global catalog for the email-based lookups of names, email addresses, and other mail-related attributes.

Because full global catalog replication can consume more bandwidth than standard domain controller replication, it is important to design a site structure to reflect the available WAN link capacity. If a sufficient amount of capacity is available, a full global catalog server can be deployed. If, however, capacity is limited, universal group membership caching can be enabled to reduce the bandwidth load.

### Understanding Multiple Forests Design Concepts Using Microsoft Identity Integration Server (MIIS) 2003

Microsoft Identity Integration Server 2003 enables out-of-the-box replication of objects between two separate AD forests. This concept becomes important for organizations with multiple Exchange implementations that want a common Global Address List for the company. Previous iterations of MIIS required an in-depth knowledge of scripting to be able to synchronize objects between two forests. MIIS 2003, on the other hand, includes built-in scripts that can establish replication between two Exchange Server 2007 AD forests, making integration between forests easier.

**NOTE**

The built-in scripts in MIIS 2003 enable synchronization only between two forests that have a full Exchange Server 2007 or Exchange Server 2003 schema. In other words, if synchronization between an Exchange 2000 forest or an Exchange 5.5 directory is required, customized scripts must be developed.

## Determining Exchange Server 2007 Placement

Previous versions of Exchange essentially forced many organizations into deploying servers in sites with greater than a dozen or so users. With the concept of site consolidation in Exchange Server 2007, however, smaller numbers of Exchange servers can service clients in multiple locations, even if they are separated by slow WAN links. For small and medium-sized organizations, this essentially means that one or two servers should suffice for the needs of the organization, with few exceptions. Larger organizations require a

larger number of Exchange servers, depending on the number of sites and users. Designing Exchange Server 2007 placement must take into account both administrative group and routing group structure. In addition, Exchange Server 2007 introduces new server role concepts, which should be understood so that the right server can be deployed in the right location.

## Understanding Exchange Server 2007 Server Roles

Exchange Server 2007 introduced the concept of server roles to Exchange terminology. In the past, server functionality was loosely termed, such as referring to an Exchange server as an OWA or front-end server, bridgehead server, or a Mailbox or back-end server. In reality, there was no *set* terminology that was used for Exchange server roles. Exchange Server 2007, on the other hand, distinctly defines specific roles that a server can hold. Multiple roles can reside on a single server, or multiple servers can have the same role. By standardizing on these roles, it becomes easier to design an Exchange environment by designating specific roles for servers in specific locations.



The server roles included in Exchange Server 2007 include the following:

- **Client access server (CAS)**—The CAS role allows for client connections via nonstandard methods such as Outlook Web Access (OWA), Exchange ActiveSync, Post Office Protocol 3 (POP3), and Internet Message Access Protocol (IMAP). CAS servers are the replacement for Exchange 2000/2003 front-end servers and can be load balanced for redundancy purposes. As with the other server roles, the CAS role can coexist with other roles for smaller organizations with a single server, for example.
- **Edge Transport server**—The Edge Transport server role is unique to Exchange 2007, and consists of a standalone server that typically resides in the demilitarized zone (DMZ) of a firewall. This server filters inbound SMTP mail traffic from the Internet for viruses and spam, and then forwards it to internal Hub Transport servers. Edge Transport servers keep a local AD Application Mode (ADAM) instance that is synchronized with the internal AD structure via a mechanism called EdgeSync. This helps to reduce the surface attack area of Exchange.
- **Hub Transport server**—The Hub Transport server role acts as a mail bridgehead for mail sent between servers in one AD site and mail sent to other AD sites. There needs to be at least one Hub Transport server within an AD site that contains a server with the Mailbox role, but there can also be multiple Hub Transport servers to provide for redundancy and load balancing.
- **Mailbox server**—The Mailbox server role is intuitive; it acts as the storehouse for mail data in users' mailboxes and down-level public folders if required. It also directly interacts with Outlook MAPI traffic. All other access methods are proxied through the CAS servers.
- **Unified Messaging server**—The Unified Messaging server role is new in Exchange 2007 and allows a user's Inbox to be used for voice messaging and fax capabilities.

Any or all of these roles can be installed on a single server or on multiple servers. For smaller organizations, a single server holding all Exchange roles is sufficient. For larger organizations, a more complex configuration might be required. For more information on designing large and complex Exchange implementations, see Chapter 4.

## Understanding Environment Sizing Considerations

In some cases with very small organizations, the number of users is small enough to warrant the installation of all AD and Exchange Server 2007 components on a single server. This scenario is possible, as long as all necessary components—DNS, a global catalog domain controller, and Exchange Server 2007—are installed on the same hardware. In general, however, it is best to separate AD and Exchange onto separate hardware wherever possible.

## Identifying Client Access Points

At its core, Exchange Server 2007 essentially acts as a storehouse for mailbox data. Access to the mail within the mailboxes can take place through multiple means, some of which might be required by specific services or applications in the environment. A good understanding of what these services are and if and how your design should support them is warranted.

### Outlining MAPI Client Access with Outlook 2007

The "heavy" client of Outlook, Outlook 2007, has gone through a significant number of changes, both to the look and feel of the application, and to the back-end mail functionality. The look and feel has been streamlined based on Microsoft research and customer feedback. Users of Outlook 2003 might be familiar with most of the layout, whereas users of Outlook 2000 and previous versions might take some getting used to the layout and configuration.

On the back end, Outlook 2007 improves the MAPI compression that takes place between an Exchange Server 2007 system and the Outlook 2007 client. The increased compression helps reduce network traffic and improve the overall speed of communications between client and server.

In addition to MAPI compression, Outlook 2007 expands upon the Outlook 2003 ability to run in cached mode, which automatically detects slow connections between client and server and adjusts Outlook functionality to match the speed of the link. When a slow link is detected, Outlook can be configured to download only email header information. When emails are opened, the entire email is downloaded, including attachments if necessary. This drastically reduces the amount of bits across the wire that is sent because only those emails that are required are sent across the connection.

The Outlook 2007 client is the most effective and full-functioning client for users who are physically located close to an Exchange server. With the enhancements in cached mode functionality, however, Outlook 2007 can also be effectively used in remote locations. When making the decision about which client to deploy as part of a design, you should keep these concepts in mind.

#### Accessing Exchange with Outlook Web Access (OWA)

The Outlook Web Access (OWA) client in Exchange Server 2007 has been enhanced and optimized for performance and usability. There is now very little difference between the full function client and OWA. With this in mind, OWA is now an even more efficient client for remote access to the Exchange server. The one major piece of functionality that OWA does not have, but the full Outlook 2007 client does, is offline mail access support. If this is required, the full client should be deployed.

#### Using Exchange ActiveSync (EAS)

Exchange ActiveSync (EAS) support in Exchange Server 2007 allows a mobile client, such as a Pocket PC device, to synchronize with the Exchange server, allowing for access to email from a handheld device. EAS also supports Direct Push technology, which allows for instantaneous email delivery to handheld devices running Windows Mobile 5.0 and the Messaging Security and Feature Pack (MSFP).



#### Understanding the Simple Mail Transport Protocol (SMTP)

The Simple Mail Transfer Protocol (SMTP) is an industry-standard protocol that is widely used across the Internet for mail delivery. SMTP is built in to Exchange servers and is used by Exchange systems for relaying mail messages from one system to another, which is similar to the way that mail is relayed across SMTP servers on the Internet. Exchange is dependent on SMTP for mail delivery and uses it for internal and external mail access.

By default, Exchange Server 2007 uses DNS to route messages destined for the Internet out of the Exchange topology. If, however, a user wants to forward messages to a smarthost before they are transmitted to the Internet, an SMTP connector can be manually set up to enable mail relay out of the Exchange system. SMTP connectors also reduce the risk and load on an Exchange server by off-loading the DNS lookup tasks to the SMTP smarthost. SMTP connectors can be specifically designed in an environment for this type of functionality.

#### Using Outlook Anywhere (Previously Known as RPC over HTTP)

One very effective and improved client access method to Exchange Server 2007 is known as Outlook Anywhere. This technology was previously referred to as RPC over HTTP(s) or Outlook over HTTP(s). This technology enables standard Outlook 2007 access across firewalls. The Outlook 2007 client encapsulates Outlook RPC packets into HTTP or HTTPS packets and sends them across standard web ports (80 and 443), where they are then extracted by the Exchange Server 2007 system. This technology enables Outlook to communicate using its standard RPC protocol, but across firewalls and routers that normally do not allow RPC traffic. The potential uses of this protocol are significant because many situations do not require the use of cumbersome VPN clients.

## Configuring Exchange Server 2007 for Maximum Performance and Reliability

After decisions have been made about AD design, Exchange server placement, and client access, optimization of the Exchange server itself helps ensure efficiency, reliability, and security for the messaging platform.

## Designing an Optimal Operating System Configuration for Exchange

As previously mentioned, Exchange Server 2007 only operates on the Windows Server 2003 operating system, and is scheduled to be able to run on the next version of the Windows Server operating system, currently referred to as Windows Longhorn. The enhancements to the operating system, especially in regard to security, make Windows Server 2003 the optimal choice for Exchange. Unless clustering (including Cluster Continuous Replication) is required, which is not common for smaller organizations, the Standard Edition of Windows Server 2003 can be installed as the OS.

**NOTE**

Contrary to popular misconception, the Enterprise Edition of Exchange can be installed on the Standard Edition of the operating system, and vice versa. Although there has been a lot of confusion on this concept, both versions of Exchange were designed to interoperate with either version of Windows.

## Avoiding Virtual Memory Fragmentation Issues

The previous iterations of Windows Server  have suffered from a problem with virtual memory (VM) fragmentation. The problem would manifest itself on systems with greater than 1GB of RAM, which run memory-intensive applications such as SQL Server or Exchange. The Advanced Server Edition of Windows 2000 Server enabled a workaround for this problem, in the form of a memory allocation switch that allocated additional memory for the user kernel.

Windows Server 2003 includes the capability of using this memory optimization technique in both the Standard and the Enterprise Editions of the software, so that the switch can now be used on any Windows Server 2003 system with more than 1GB of physical RAM. The switch is added to the end of the `boot.ini` file.

The `/3GB` switch tells Windows to allocate 3GB of memory for the user kernel, and the `/USERVA=3030` switch optimizes the memory configuration, based on tests performed by Microsoft that determined the perfect number of megabytes to allocate for optimal performance and the least likely instance of VM fragmentation. This setting only applies to the 32-bit version of Windows 2003, so it would not apply to Exchange 2007 servers but would apply to 32-bit domain controllers and any other supporting 32-bit servers in an Exchange 2007 environment.

## Configuring Disk Options for Performance

The single most important design element, which improves the efficiency and speed of Exchange, is the separation of the Exchange database and the Exchange logs onto a separate hard drive volume. Because of the inherent differences in the type of hard drive operations performed (logs perform primarily write operations, databases primarily read),

separating these elements onto separate volumes dramatically increases server performance. Keep these components separate in even the smallest Exchange server implementations. Figure 3.3 illustrates some examples of how the database and log volumes can be configured.

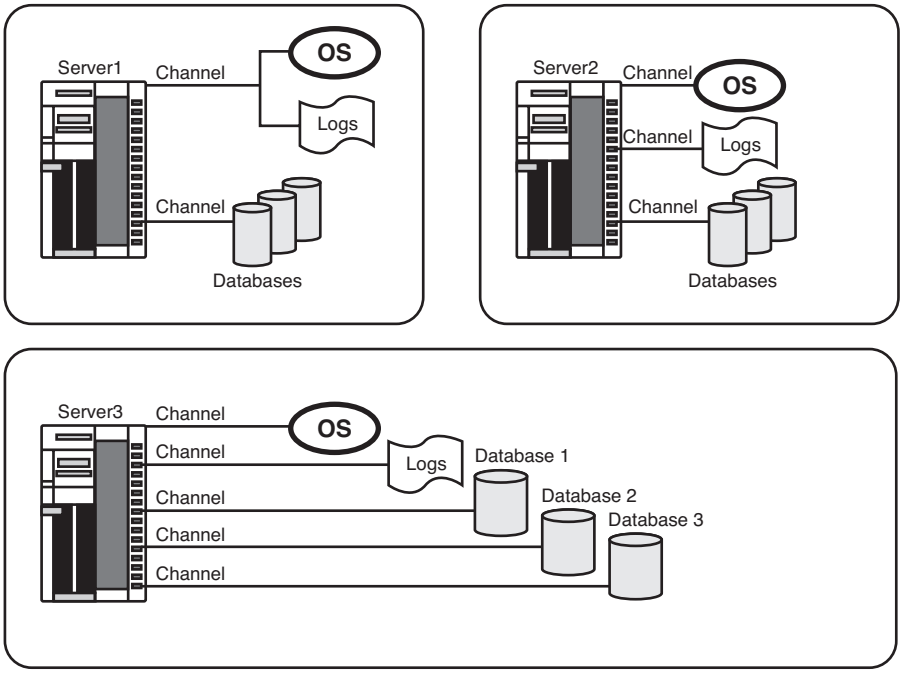


FIGURE 3.3 Database and log volume configuration.

On Server1, the OS and logs are located on the same mirrored C:\ volume and the database is located on a separate RAID-5 drive set. With Server2, the configuration is taken up a notch, with the OS only on C:\, the logs on D:\, and the database on the RAID-5 E:\ volume. Finally, Server3 is configured in the optimal configuration, with separate volumes for each database and a volume for the log files. The more advanced a configuration, the more detailed and complex the drive configuration can get. However, the most important factor that must be remembered is to separate the Exchange database from the logs wherever possible.

## Working with Multiple Exchange Databases and Storage Groups

The Enterprise Edition of Exchange Server 2007 not only enables databases of larger than 75GB, it also enables the creation of multiple separate databases on a single server. This concept gives great flexibility in design while enabling reduced downtime and increased performance.

A storage group is a logical grouping of databases that share a single set of logs. Each Exchange Server 2007 Enterprise system can handle a maximum of 50 storage groups per server. Each storage group can contain a maximum of five databases each, although the total number of databases on a server cannot equal more than 50.

**NOTE**

If Cluster Continuous Replication (CCR) is to be used, it is important to note that CCR only supports a single database per storage group. Also, Microsoft recommends no more than 30 databases on a server running CCR.

In practice, however, each instance of a storage group that is created uses a greater amount of resources, so it is wise to create additional storage groups only if absolutely necessary. Multiple databases, on the other hand, can solve several problems:

- **Reduce database restore time**—Smaller databases take less time to restore from tape. This concept can be helpful if there is a group of users who require quicker recovery time (such as management). All mailboxes for this group could then be placed in a separate database to provide quicker recovery time in the event of a server or database failure.
- **Provide for separate mailbox limit policies**—Each database can be configured with different mailbox storage limits. For example, the standard user database could have a 200-MB limit on mailboxes, and the management database could have a 500-MB limit.
- **Mitigate risk by distributing user load**—By distributing user load across multiple databases, the risk of losing all user mail connectivity is reduced. For example, if a single database failed that contained all users, no one would be able to mail. If those users were divided across three databases, however, only one third of those users would be unable to mail in the event of a database failure.

**NOTE**

One disadvantage to multiple databases is that the concept of single-instance storage is lost across databases. Single-instance storage occurs when only one copy of an email message sent to multiple people is stored on the server, dramatically reducing the space needed to store mass mailings. Each separate database must keep a copy of mass mailings, however, which increases the aggregate total size of the databases.

## Understanding Clustering for Exchange Server 2007

Exchange Server 2007 is configured to use Windows Server 2003 clustering for enhanced redundancy and increased uptime. Clustering is an expensive option, but one that will increase reliability of the Exchange Server 2007 implementation.

Clustering options with Exchange Server 2007 have significantly changed over those available in previous versions. Traditional, shared storage clustering is now referred to as a Single Copy Cluster. New options for clustering databases across geographical locations automatically using asynchronous synchronization of log files is now available and is referred to as Cluster Continuous Replication (CCR). More information on clustering with Exchange 2007 can be found in Chapters 4, "Architecting an Enterprise-Level Exchange Environment," and 31, "Continuous Backups, Clustering, and Network Load Balancing in Exchange Server 2007."

**NOTE**

Microsoft no longer supports a full active-active clustering configuration. Consequently, at least one cluster node should be configured as passive. With eight-way clustering, for example, this means that seven nodes can be active and one node must be passive.



### Monitoring Design Concepts with Microsoft Operations Manager 2005

The enhancements to Exchange Server 2007 do not stop with the improvements to the product itself. New functionality has been added to the Exchange Management Pack for Microsoft Operations Manager (MOM) that enables MOM to monitor Exchange servers for critical events and performance data. The MOM Management Pack is preconfigured to monitor for Exchange-specific information and to enable administrators to proactively monitor Exchange servers. For more information on using MOM to monitor Exchange Server 2007, see Chapter 20, "Using Microsoft Operations Manager to Monitor Exchange Server 2007."

## Securing and Maintaining an Exchange Server 2007 Implementation

One of the greatest advantages of Exchange Server 2007 is its emphasis on security. Along with Windows Server 2003, Exchange Server 2007 was developed during and after the Microsoft Trustworthy Computing initiative, which effectively put a greater emphasis on security over new features in the products. In Exchange Server 2007, this means that the OS and the application were designed with services "Secure by Default."

With Secure by Default, all nonessential functionality in Exchange must be turned on if needed. This is a complete change from the previous Microsoft model, which had all services, add-ons, and options turned on and running at all times, presenting much larger security vulnerabilities than was necessary. Designing security effectively becomes much easier in Exchange Server 2007 because it now becomes necessary only to identify components to turn on, as opposed to identifying everything that needs to be turned off.

In addition to being secure by default, Exchange Server 2007 server roles are built in to templates used by the Security Configuration Wizard (SCW), which was introduced in

Service Pack 1 for Windows Server 2003. Using the SCW against Exchange Server helps to reduce the surface attack area of a server.

### Patching the Operating System Using Windows Software Update Services

Although Windows Server 2003 presents a much smaller target for hackers, viruses, and exploits by virtue of the Secure by Default concept, it is still important to keep the OS up to date against critical security patches and updates. Currently, two approaches can be used to automate the installation of server patches. The first method involves configuring the Windows Server 2003 Automatic Updates client to download patches from Microsoft and install them on a schedule. The second option is to set up an internal server to coordinate patch distribution and management. The solution that Microsoft supplies for this functionality is known as Windows Software Update Services (WSUS).

WSUS enables a centralized server to hold copies of OS patches for distribution to clients on a preset schedule. WSUS can be used to automate the distribution of patches to Exchange Server 2007 servers, so that the OS components will remain secure between service packs. WSUS might not be necessary in smaller environments, but can be considered in medium-sized to large organizations that want greater control over their patch management strategy.

### Implementing Maintenance Schedules

Exchange still uses the Microsoft JET Database structure, which is effectively the same database engine that has been used with Exchange from the beginning. This type of database is useful for storing the type of unstructured data that email normally carries, and has proven to be a good fit for Exchange Server. Along with this type of database, however, comes the responsibility to run regular, scheduled maintenance on the Exchange databases on a regular basis.

Although online maintenance is performed every night, it is recommended that Exchange databases be brought offline on a quarterly or, at most, semiannual basis for offline maintenance. Exchange database maintenance utilities, `eseutil` and `isinteg`, should be used to compact and defragment the databases, which can then be mounted again in the environment.

Exchange databases that do not have this type of maintenance performed run the risk of becoming corrupt in the long term, and will also never be able to be reduced in size. Consequently, it is important to include database maintenance into a design plan to ensure data integrity.

## Summary

Exchange Server 2007 offers a broad range of functionality and improvements to messaging and is well suited for organizations of any size. With proper thought for the major design topics, a robust and reliable Exchange email solution can be put into place that will perfectly complement the needs of any organization.

When Exchange design concepts have been fully understood, the task of designing the Exchange Server 2007 infrastructure can take place.

## Best Practices

The following are best practices from this chapter:

- Use site consolidation strategies to reduce the number of Exchange servers to deploy.
- Separate the Exchange log and database files onto separate physical volumes whenever possible.
- Install Exchange Server 2007 on Windows Server 2003 R2 Edition when possible.
- Integrate an antivirus and backup strategy into Exchange Server design.
- Keep a local copy of the global catalog close to any Exchange servers.
- Implement quarterly or semiannual maintenance procedures against Exchange databases by using the `isinteg` and `eseutil` utilities.
- Keep the OS and Exchange up to date through service packs and software patches, either manually or via Windows Software Update Services.
- Keep the AD design simple, with a single forest and single domain, unless a specific need exists to create more complexity.
- Identify the client access methods that will be supported and match them with the appropriate Exchange Server 2007 technology.
- Implement DNS in the environment on the AD domain controllers.