

realtimepublishers.com<sup>tm</sup>

*Tips and Tricks*  
*Guide<sup>tm</sup> To*

**Secure  
Messaging**

*Jim McBee*

**Note to Reader:** This book presents tips and tricks for six email security topics. For ease of use, the questions and their solutions are divided into topics, and each question is numbered based on the topic, including

- Topic 1: Strategies for Defending Email Infrastructure
- Topic 2: Policies and Procedures
- Topic 3: Architecture and Deployment Considerations
- Topic 4: Antivirus and Anti-Spam Strategies and Best Practices
- Topic 5: Firewall Strategies and Best Practices
- Topic 6: Protecting and Controlling Sensitive Information in Email

Topic 1: Strategies for Defending Email Infrastructure .....	1
Q 1.4: How do you lock down an Exchange Server? .....	1
Disabling Services .....	1
Disabling Mailbox Server Services .....	2
Disabling Front-End Server Services.....	4
Internet Exposure Protection and Redundancy .....	5
Applying Limits .....	7
Mailbox Limits.....	7
Message Size Limits .....	8
Maximum Store Size.....	9
Avoiding Unnecessary Exposure.....	11
Topic 2: Policies and Procedures.....	12
Q 2.4: How do I access users' mailboxes to extract a dangerous message?.....	12
Mailbox Access Permissions .....	13
Running ExMerge.....	16
Topic 3: Architecture and Deployment Considerations.....	19
Q 3.4: What are best practices to follow when deploying servers? .....	19
Building a Solid Server Platform.....	19
Checking Your Work.....	21
Infrastructure Protection .....	24
Bad Practices.....	24
Topic 4: Protecting and Controlling Sensitive Information in Email.....	25
Q 4.4: What options are available for antivirus and spam protection?.....	25

---

Managed Providers .....	25
Protection in the DMZ .....	26
Protection on the Mail Server .....	26
Topic 5: Firewall Strategies and Best Practices.....	27
Q 5.4: What are best practices for configuring a firewall to support better email security? .....	27
Topic 6: Protecting and Controlling Sensitive Information in Email .....	28
Q 6.4: How does Enterprise Rights Management work?.....	28
Publishing Licenses and Usage Licenses.....	28
RMS .....	29
Publisher Setup .....	30
RMS-Enabled Applications .....	30
Rights Management Client Software .....	30
Using RMS for the First Time .....	31
Creating Protected Content .....	32
Consuming Protected Content .....	35

## **Copyright Statement**

© 2006 Microsoft Corporation. All rights reserved.

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library. All leading technology guides from Realtimepublishers can be found at <http://nexus.realtimepublishers.com>.]


## **Topic 1: Strategies for Defending Email Infrastructure**

### **Q 1.4: How do you lock down an Exchange Server?**

**A:** Most modern mail systems are fairly robust and can resist a lightweight attempt at a Denial of Service (DoS) attack. However, a determined individual or group of individuals can still wreak havoc on an unprotected or unprepared mail system if they approach their attack from the right angle. There are a few simple steps that organizations can take to reduce the likelihood that their publicly exposed mail services are compromised.

These steps include applying the appropriate messaging system limits to the server so that it cannot exceed the capacity that you have planned for each user's mailbox. Furthermore, you can lock down the users' ability to automatically forward information outside of the organization.

Reducing the number of services that a server is running or reducing the number of ports that are opened can also help reduce the possibility of DoS attacks by decreasing the server's surface area. If services are exposed to the Internet, place proxy or relay systems between the Internet and the internal mail servers.

 Mail servers operate best and most securely when they have only the necessary applications and services to perform the necessary functions. Avoid installing too many services or third-party applications if they are not necessary for the server's designated role.

### ***Disabling Services***

One of the first steps to better protection for any Windows server is to reduce the number of services that the server is running and reduce the number of TCP or UDP ports that are open. Although specific services or ports may not currently be employed by attackers to compromise a system, there is no guarantee that weaknesses in those services or ports won't be discovered in the future.

The different services that you chose to remove or disable on any Exchange Server will vary based on your organization's needs. The services you can disable will also vary based on the role that particular server plays in the organization.

## Disabling Mailbox Server Services

Exchange Server systems that function as a mailbox or public folder server (a back-end server) have a different set of services that may be disabled than servers that operate as front-end servers. The services that exist on a Windows Server 2003 (WS2K3) machine will vary based on the additional Windows or third-party components that have been installed. Table 1.5 shows a list of the services that may be disabled on a mailbox server and under what conditions.

Service	Function
Alerter	The Alerter service is disabled by default in WS2K3. It is used by the operating system (OS) and some services to send network administrative alerts.
Application Experience Lookup Service	Application Experience Lookup Service handles lookup requests for application compatibility. This functionality is usually not necessary on a server.
Application Layer Gateway Service	Application Layer Gateway Service handles protocol support for plug-ins that use Internet Connection Sharing or the Windows Firewall. This functionality is usually not necessary on a server.
Computer Browser	The Computer Browser service provides the Network Neighborhood function and ensures that this computer is listed in the Network Neighborhood. Exchange does not use this service and the clients do not depend on it. However, some third-party applications use this service to locate computers on the network, so consider this possibility when disabling the service.
Error Reporting Service	Error Reporting Service is responsible for collecting and reporting information about application crashes to Microsoft or to an internal error reporting service. If you do not require this service, it can be set to Manual.
Intersite Messaging	The Intersite Messaging service works with Active Directory (AD) domain controllers to exchange AD replication messages and site routing information to be transferred. If the machine is not functioning as a domain controller, this service is not necessary.
Messenger	The Messenger service is responsible for receiving network pop-up alerts and displaying them on the computer's console.
Microsoft Exchange Event	The Microsoft Exchange Event service monitors and runs Exchange 5.5-compatible folder event scripts. This service is not necessary if you have no such event scripts registered.

Microsoft Exchange IMAP4	The IMAP4 service is disabled by default on a freshly built Exchange 2003 server but may remain enabled if the server was upgraded from Exchange 2000. If you have no IMAP4 clients, this service can be disabled. Doing so will close TCP port 143.
Microsoft Exchange Management	The Exchange Management service provides an interface between Exchange functions and Windows Management Instrumentation (WMI). If this service is disabled, WMI management functions (such as monitoring) will not work and the Exchange Message Tracking center will not be able to look up data in the server's message tracking logs.
Microsoft Exchange MTA Stacks	The Exchange MTA is responsible for delivering mail to Exchange 5.5 servers that are in the same Exchange site or that connect to this server through an Exchange 5.5 Site Connector. The MTA is also used for some types of gateway software packages and is used if you have X.400 connectors. Disabling this service will close TCP port 102.
Microsoft Exchange POP3	The POP3 service is disabled by default on a freshly built Exchange 2003 server but may remain enabled if the server was upgraded from Exchange 2000. If you have no POP3 clients, this service can be disabled. Doing so will close TCP port 110.
Microsoft Exchange Site Replication Service (SRS)	The Exchange SRS emulates an Exchange 5.5 directory service for the purposes of serving as a go-between for the AD and Exchange 5.5 servers. If there are not Exchange 5.5 servers in your organization, this service should be disabled. Disabling this port will close TCP port 379 and a dynamic RPC port above 1024.
Microsoft Search	The Microsoft Search component is used to create full-text indexes of public folder or mailbox stores. If this feature is not being used, you can stop this service.
Network News Transport Protocol (NNTP)	The NNTP service is used to synchronize Exchange public folders or newsgroups with USENET NNTP servers. It is also used by NNTP clients to access Exchange public folders or newsgroups. This service is disabled by default on a fresh Exchange 2003 installation, but may have been enabled in Exchange 2000 and remain enabled after an upgrade.
Task Scheduler	The Task Scheduler service provides a method of scheduling scripts and programs to run. Exchange does not require this service, but some backup software packages and other third-party applications do. If the service is not required, it can be disabled.

Telnet	The Telnet service is disabled by default but many administrators enable it. It provides a command-line interface for server management but offers no encryption of the data crossing the network between the Telnet client and the service. As long as it is closed, TCP port 23 is not available.
WinHTTP Web Proxy Auto-Discovery Service	WinHTTP Web Proxy Auto-Discovery is used to discover the Web proxy configuration. This functionality is not normally necessary on a server.

**Table 1.5: Windows and Exchange services that can potentially be disabled on a mail server.**

Table 1.5 is not an all-inclusive list of services that you may find running on an Exchange 2003 server, but it includes a list of common services that may be enabled by default or because of an upgrade from a previous version of Exchange.

Prior to disabling any of these services, confirm that the service is not required by your organization or an application running on the server. When making changes in your organization, such as installing new applications or supporting new features, keep in mind that these services may have been disabled but are now required.

## Disabling Front-End Server Services

Exchange front-end servers provide different types of remote client access. A front-end server may be dedicated to HTTP client access functionality—such as Outlook Web Access (OWA), Windows Mobile ActiveSync, or HTTP over RPC proxy functions. A front-end server can also function as a Simple Mail Transfer Protocol (SMTP) or X.400 bridgehead server for handling messages in and out of the routing group or to an external organization. Depending on the size of the organization, a single server (or load-balanced cluster) may perform all these functions, or the client access functions may be split off from the messaging bridgehead functions.



The list of services provided in Table 1.5 also applies to front-end servers; however, there are a few more services that may be disabled depending on the server's role and functions. Table 1.6 shows a list of additional services that you may be able to disable if they are unnecessary.


Service	Function
SMTP	The SMTP service is responsible for delivering SMTP mail between Exchange 2003 servers and other email servers on the Internet. This service must remain running on all Exchange 2003 servers that host mailboxes or act as a messaging bridgehead of any type. However, on a front-end server dedicated to just client access, it can be disabled. When disabled, TCP port 25 is no longer accessible.
Microsoft Exchange Information Store	The Information Store service is responsible for managing the mailbox and public folder stores, running the database engine, and providing client access via MAPI or Internet protocols. The Information Store service and the default mailbox store are necessary on bridgehead servers, but if a server is dedicated to just client access, the service can be disabled. If disabled, a random RPC port above 1024 will be closed.
Microsoft Exchange Routing Engine	The routing engine service handles the sharing of link state information between Exchange 2000/2003, servers which, in turn, provides an optimal message routing topology. If the Exchange front-end server is functioning only as a client-access server, the routing engine can be disabled. Doing so will close TCP port 691.

**Table 1.6: Windows and Exchange services that can potentially be disabled on a front-end server.**

### ***Internet Exposure Protection and Redundancy***

A determined hacker or group of hackers can launch a DoS attack against your email servers by doing something as simple as sending many simultaneous inbound requests via an open Internet port. Most organizations open Web services (usually HTTPS using port 443) and, of course, SMTP (port 25).


Sometimes, slightly evil deeds bring unintended consequences, as in the unholy union between hackers and spammers. Hackers have built large, distributed spamming systems using virus and worms. A single SMTP server can have dozens or hundreds of inbound SMTP connections all trying to deliver spam using dictionary message delivery techniques. All these sessions come from different IP addresses, thus, blocking them is difficult.

 An important line of defense against DoS attacks and unwanted intrusions is to ensure that your perimeter is correctly sealed and protected, including a properly configured firewall. Audits should be conducted regularly to ensure that no inbound port is left open if it is no longer necessary.

For small and midsized organizations, inbound HTTPS or SMTP might actually be routed directly to mailbox servers. A DoS attack against your organization could cripple not only inbound email or OWA clients but also mail service for the entire organization.

Protecting the mailbox servers from this type of attack becomes an important part of your security defense. The most effective method of stopping a direct attack on your internal mailbox servers is to simply not expose them to the Internet. Instead, use reverse proxy solutions for inbound HTTPS and SMTP relays or third-party inspection systems to protect SMTP.

For scalability and redundancy, additional gateways or proxies can be put in place. Figure 1.16 shows an organization that has implemented this type of defense. They have placed SMTP relays in their DMZ and ISA Servers acting as reverse proxies. In this case, each of these have implemented the network load-balancing service so that, for inbound connections, all the machines appear as one for external clients.



**Figure 1.16: Providing protection for internal mail server resources.**

## Applying Limits

Most mail administrators know that storage limits for mailboxes and message size limits for inbound and outbound mail is a good idea because these limits help in planning expected capacity and knowing how much data you need to backup and restore. However, applying limits also has security implications.

An intruder intent on shutting down an organization's mail servers might start sending many large messages to multiple mailboxes within the organization. With as little as a few hours, a mail server's transaction log disk, database disk, or message queue folder may overflow due to the larger than expected messaging load. For organizations whose management is reluctant to impose limitations on their users, this threat may give you the additional ammunition necessary to convince your boss that you need mailbox or message size limits.

## Mailbox Limits

Mailbox limits allow you to impose three size limits to a mailbox. Figure 1.17 shows the Limits property page on an Exchange mailbox store. Messaging limits can be configured directly on the Limits property page of an Exchange mailbox store or on a selected group of mailbox stores using an Exchange Server mailbox store policy. If you have many mailbox stores, using a mailbox store policy is more efficient.

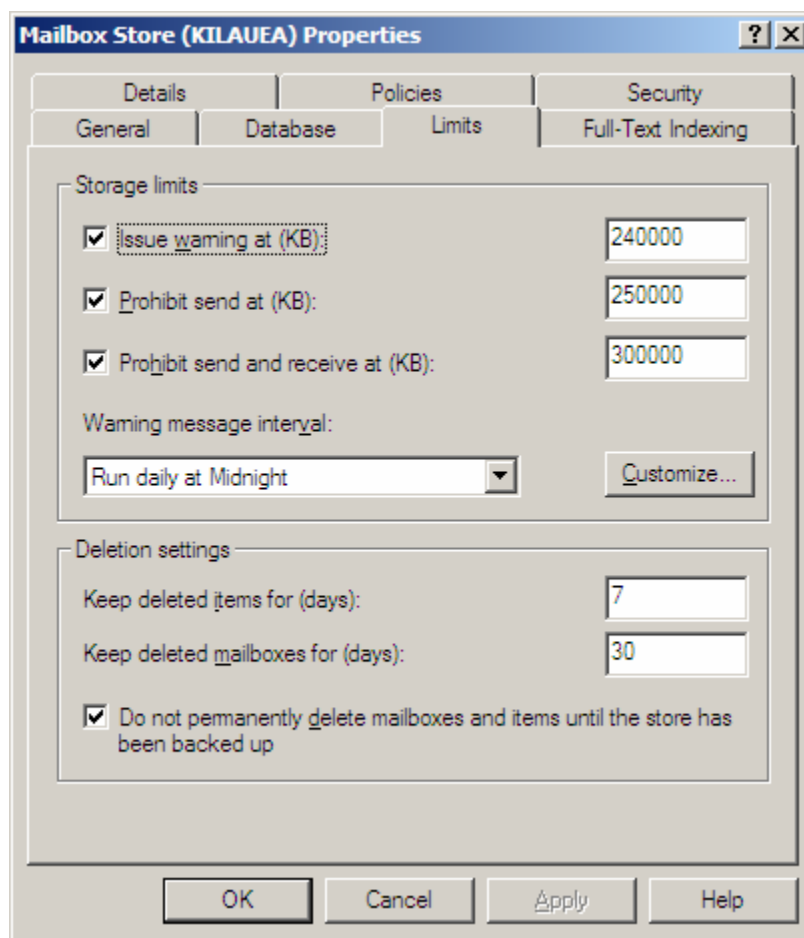


Figure 1.17: Applying storage limits.

The default mailbox limits can be overridden on a user-by-user basis by locating the user in Active Directory Users and Computers, displaying the user's properties, going to the Exchange General property page, and clicking Storage Limits.

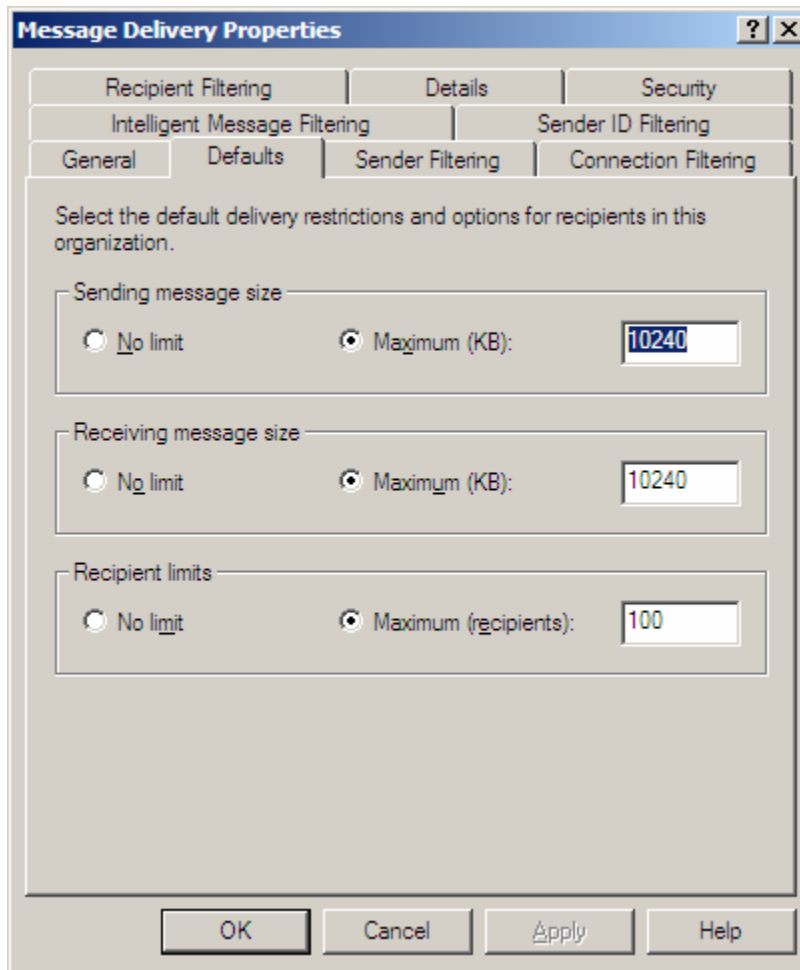
I recommend that all three limits be applied; the actual limits used vary wildly from organization to organization based on disk and backup capacity as well as the amount of data that users need to keep in their mailboxes. If you are concerned about a DoS attack that fills up a mailbox, the Prohibit Send and Receive limit will be useful to you. Unfortunately, just setting this limit does not give users any warning that their mailboxes are going to back up. An example scenario for mailbox limits would include publishing a maximum mailbox size limit of 250MB but configuring the size limits as such:

- Warning limit: 240MB
- Prohibit send at: 250MB
- Prohibit send and receive at: 300MB

Users will start seeing a daily warning message when their mailboxes hit 240MB, but at 250MB, they will no longer be able to send messages using MAPI or OWA clients. When the Prohibit Send and Receive limit is reached, the mailbox closes and will not accept any more mail, so you want to make sure that you give your users plenty of warning.

### **Message Size Limits**

Earlier versions of Exchange did not apply a maximum inbound and outbound message size limit. With Exchange 2003, there is a global limit applied that cannot be exceeded by users. This limit is configured under the Global Settings on the Message Delivery properties (the default property page is shown in Figure 1.18).



**Figure 1.18: Global message size limits.**

For most organizations, the 10MB maximum message size is a reasonable limit, but each organization should evaluate this value to determine whether they typically send and receive larger message sizes or if 10MB is too large.

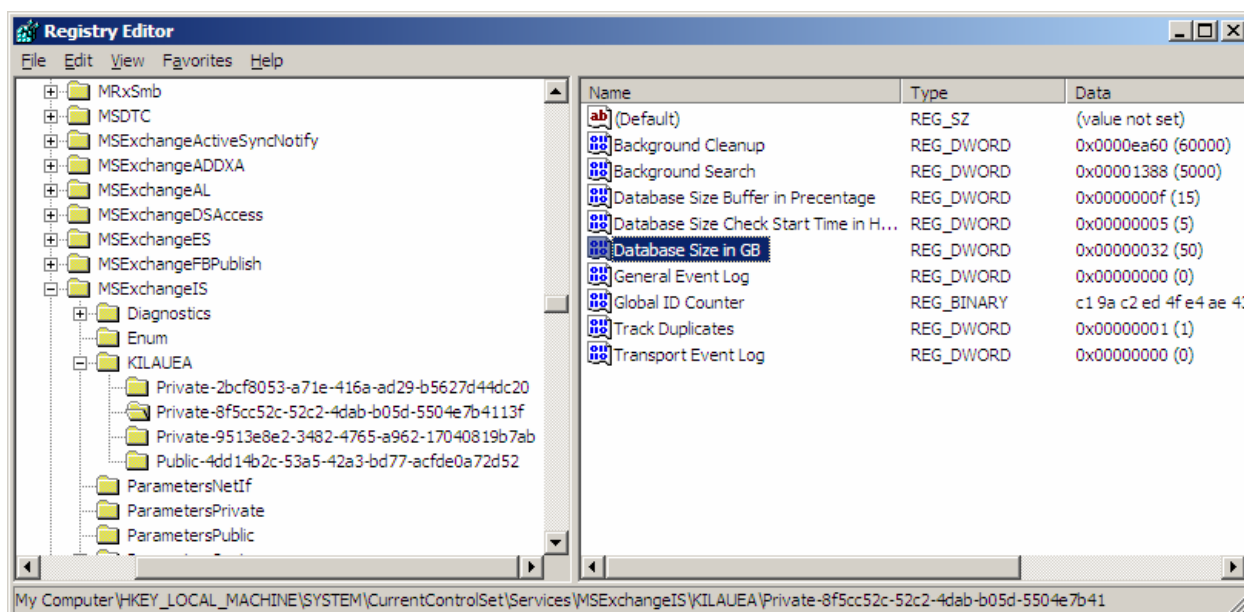
### Maximum Store Size

Exchange 2003 Service Pack 2 (SP2) introduced a feature that generated a lot of excitement for Exchange 2003 Standard Edition customers. SP2 increased the maximum mailbox store size from 16GB to 18GB; but the real excitement about this service pack is that using a registry value, the maximum store size can be increased to 75GB on Standard Edition.

You might wonder what this increased store size has to do with security and DoS. Well, the registry key works with Exchange 2003 Enterprise Edition as well as Standard Edition. It allows you to configure a maximum store size for Exchange 2003 Enterprise Edition. So, for example, suppose that you have five mailbox stores on the F drive and the maximum amount of disk space you have available on the F drive is 300GB. You decide that the maximum store size should be no more than 50GB (you want to leave yourself some headroom here because the maximum store size that is calculated does not include the white space or deleted items). You could configure each mailbox store to be a maximum of 50GB.

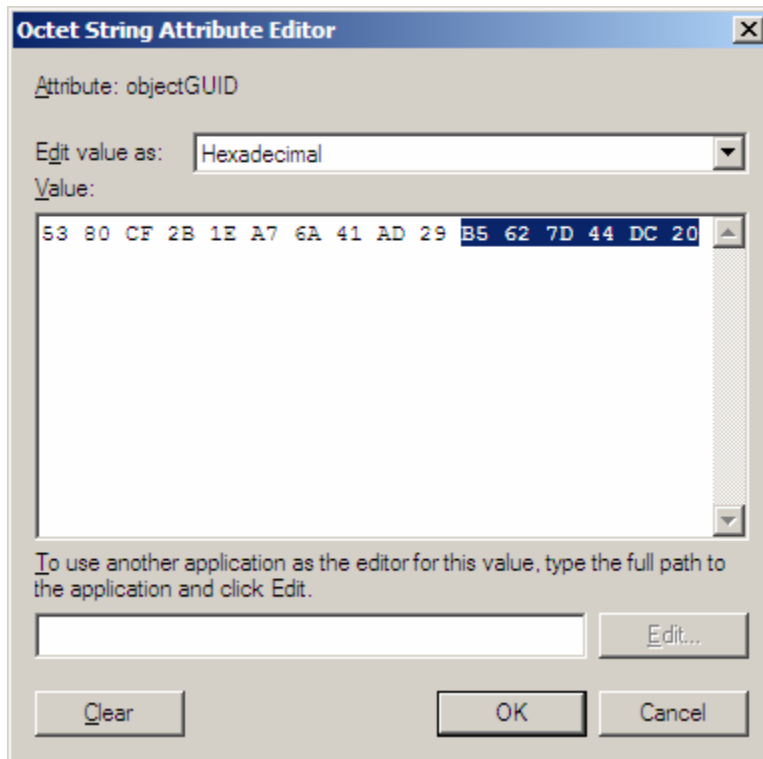
When a mailbox store size hits 50GB worth of data, it will be dismounted. Although this approach might not be ideal, it will prevent a single store from growing uncontrollably and forcing all the mailbox stores to be dismounted when the disk drive runs out of disk space.

To enable a maximum store size, you need to edit the registry and create a new registry value for each mailbox store. First, you need to locate the necessary registry key; each store has a private globally unique identifier (GUID) associated with the store name. Figure 1.19 shows an example of these registry keys. Because Figure 1.19 shows a server that is running Exchange 2003 Enterprise Edition, there is more than one mailbox store.



**Figure 1.19: Setting the maximum store size.**

You can match the GUID in the registry (well, at least part of it) with the mailbox store's objectGUID attribute by examining the mailbox store object in ADSIEdit. Take the last part of the GUID from the registry (in Figure 1.19, this value is B5627D44DC20) and match that up with the last part of the digits in the objectGUID attribute (see Figure 1.20).



**Figure 1.20: Examining a mailbox store's objectGUID attribute.**

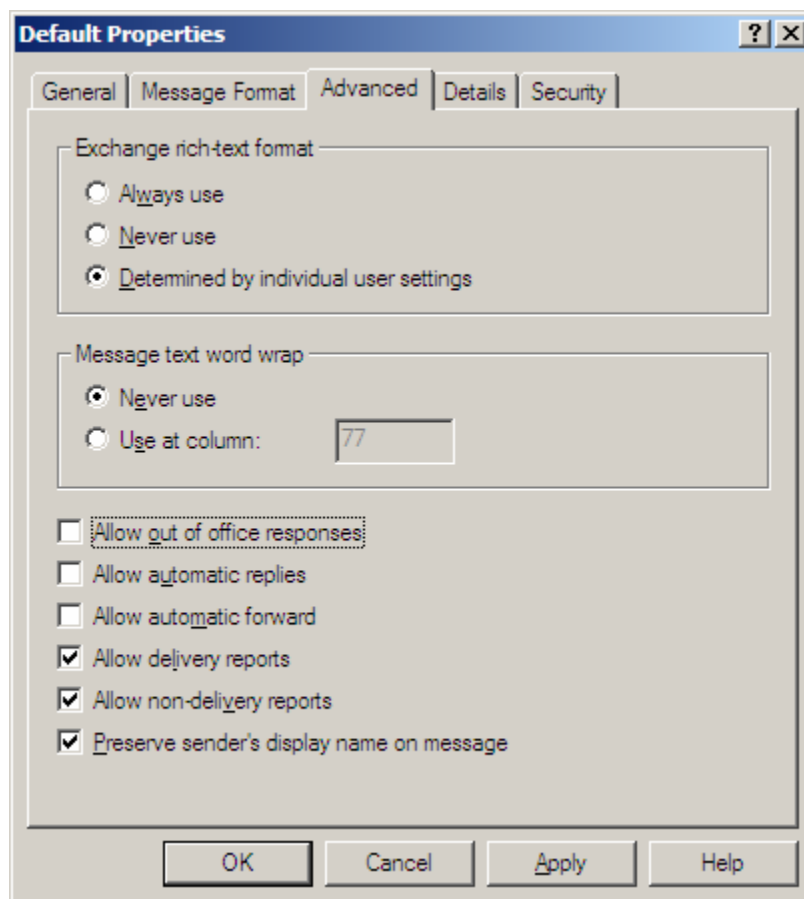
In each registry key that represents a mailbox store for which you would like to restrict the maximum store size, create a REG\_DWORD value called Database Size Limit in GB value and set it to the maximum size (don't forget to click the Decimal radio button) that you want to enable the store to grow.

### **Avoiding Unnecessary Exposure**

One of the most common ways that intruders gain access to an organization's resources is through social engineering. An intruder learns via an auto-reply message or an out-of-office message that the IT Manager Paul Agamata is out of the office. He can then use that information to attempt to gain access by posing as Paul or by asking for something on behalf of Paul.

Automatic replies and forwards can also be dangerous if a piece of sensitive data is emailed to a user, but that user has that information automatically forwarded to an external mail system. Once the data is forwarded out of the organization, there is a greater chance that it may be used by someone with less than pure intentions.

A fresh Exchange 2003 installation automatically blocks automatic replies, automatic forwards, and out of office messages, but these are frequently enabled by novice administrators. These are configured under Global Settings, Internet Message Formats. The default Internet Message Format is called Default and it applies to all outbound mail. Figure 1.21 shows the Advanced property page for this message format; you can see that automatic replies, automatic forwards, and out of office messages are not enabled.



**Figure 1.21: Defining automatic message handling default settings.**

If you have specific business partners or Internet domains to which you want to allow these automatic responses, you can configure additional Internet Message Formats for specific domains.

## **Topic 2: Policies and Procedures**

### **Q 2.4: How do I access users' mailboxes to extract a dangerous message?**

**A:** In some situations, it might be necessary to open one or more mailboxes in order to extract a sensitive message that was sent to users by accident or to remove a virus or worm that has entered many users' mailboxes. Although this task is technically straightforward, the administrative or security policy ramifications may be more complex. Any administrator that is going to perform such an operation should make sure that they are not violating a Human Resources, data retention, or other security policy.



For Exchange Server, the permission necessary to access other users' mailboxes is not given by default and requires tweaking in order to grant this permission. There is a good reason for this—the ramifications of allowing even a trusted administrator access to some or all mailboxes may be a touchy subject.

The most obvious way to remove items from a user's mailbox is to use Outlook or Outlook Web Access (OWA) to simply delete the message. If you support more than two or three mailboxes, though, this method quickly becomes a protracted and frustrating exercise.

The tool of choice for such exercises is ExMerge; the ExMerge tool can be downloaded from Microsoft's Exchange downloads page. Getting the tool and installing it (just decompress it and copy it to the \program files\exchsrvr\bin folder) is the easy part, getting the necessary permissions to access other users' mailboxes is the more difficult part.

### Mailbox Access Permissions

By default, all users or groups that have been delegated administrative permissions to the Exchange organization and administrative groups have the Receive As and Send As permissions set to Deny. This setting is available on the Security property page of any object; the entire Exchange organization's Security property page is shown in Figure 2.3.

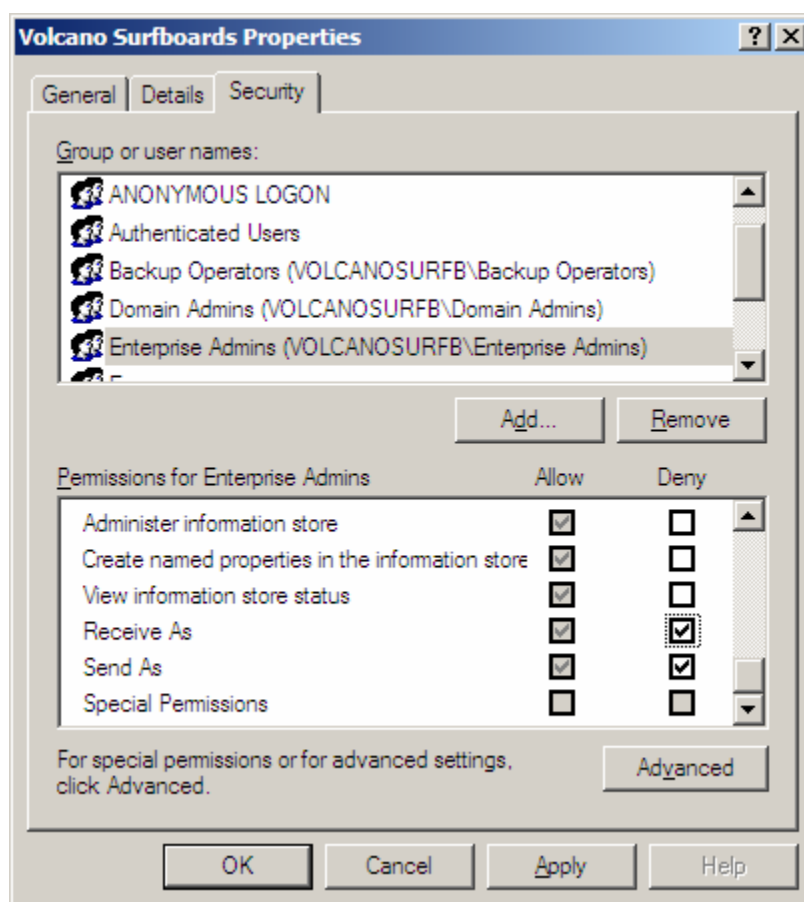


Figure 2.3: Security property page seen from Exchange System Manager.

If you are familiar with the Exchange System Manager interface for Exchange 2000 and 2003, you probably realize that the Security property page does not show up on all objects. This includes the top-level Exchange organization and the administrative groups. This is by design so that someone with Exchange Full Administrator permissions (which includes the ability to change the permissions list) cannot accidentally change the denied permissions. These permissions are explicitly set at the organizational level; the only place they can be removed is at this level.

The Enterprise Admins and Domain Admins groups both have the Send As and Receive As permissions denied. The user account that ran the initial Exchange forest preparation (forestprep) is also denied the Send As and Receive As permissions. In many cases, this is the administrator of the forest root domain.



The Send As permission allows someone to send a message as another user; this permission is used by the Exchange Server software to deliver mail. The Receive As permission allows mailboxes to be accessed.

You can enable the Security property page view in Exchange System Manager through the registry. Create a REG\_DWORD value called ShowSecurityPage in the HKEY\_CURRENT\_USER\Software\Microsoft\Exchange\ExAdmin registry key. Set this value to 1, and the Security page should show up next time you check the organization or administrative group objects using Exchange System Manager.

There are a couple of ways to approach the delegation of the necessary permissions. The first method that many novice administrators think of is to merely remove the Receive As and Send As Deny permissions at the organization level for the Enterprise Admins and the Domain Admins. This method is effective, but may yield undesirable results because this means that all members of either Domain Admins or Enterprise Admins will be able to open anyone's mailbox.

The second approach is to delegate a user (or group) permissions to the Exchange organization (or to a single administrative group) using the Delegation Wizard in Exchange System Manager. Once the permissions have been delegated, remove the Receive As and Send As Deny permissions. This option will work as well, but it gives that user or group more permissions than they really require.

I am a big proponent of accountability, segmentation of user roles, and practicing the principle of least permissions. For these reasons, I recommend the following procedure for creation of a user that has permissions only to access mail. Then carefully protect that user so that only a selected group of administrators have access to that user account. The procedure I am outlining is for an entire Exchange organization, but you can apply it to an individual administrative group just as easily as the entire organization. The user and group names are also examples, so use whatever is appropriate for your organization.

- Create a security group called Exchange Full Mailbox Access
- Create a user called ExMergeUser, set a strong password on the user, and make the user a member of the Exchange Full Mailbox Access security group. This user does not need a mailbox.
- Delegate the Exchange Full Mailbox Access security group the Exchange View Only Administrator role to the Exchange organization object. This group does not need Exchange administrator abilities.
- Using the Exchange System Manager, access the Security property page of the Exchange organization, locate the Exchange Full Mailbox Access security group in the list of groups or users, and select the Allow box for the Receive As permission (Figure 2.4) The Send As permission is not necessary and is excessive.

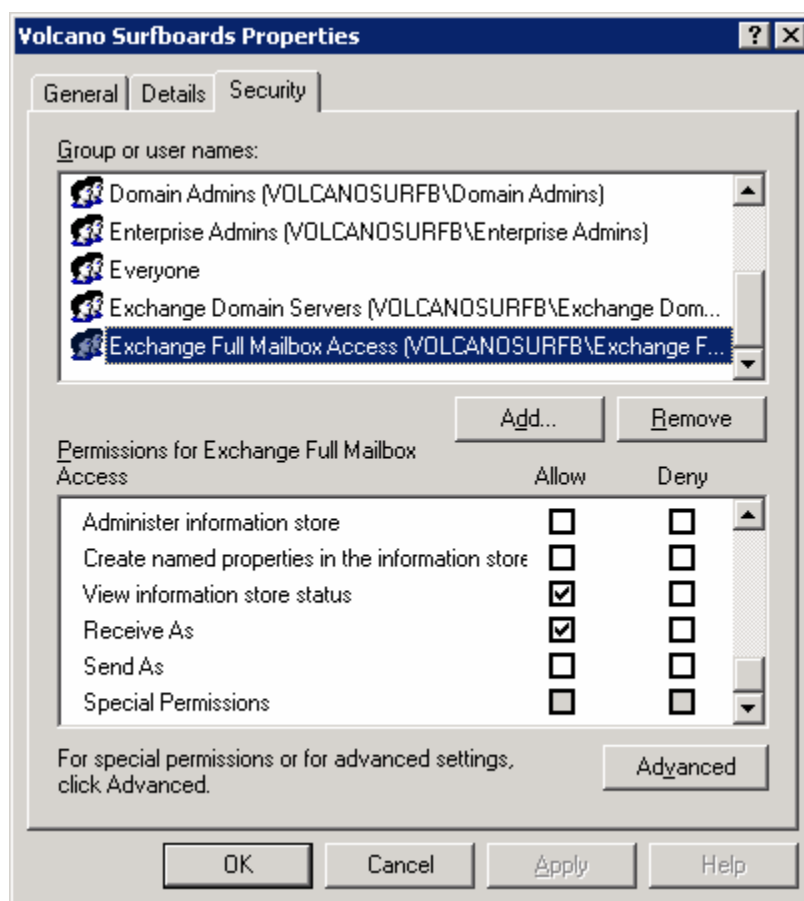


Figure 2.4: Modifying the View Only permissions to include Receive As.

You might need to wait up to 2 hours for the permissions to take effect or you can stop and restart the information store if you just can't wait.


If the ExMergeUser is a member of Domain Admins, Enterprise Admins, or the Exchange Domain Servers groups, the permissions will not be sufficient because each of these groups is explicitly denied permission to access mailboxes. If you have failures running ExMerge and the ExMerge.log file shows the following errors, you don't have the correct permissions or you have not given them time to sufficiently replicate and take effect:

[16:58:52] Error opening message store (MSEMS). Verify that the Microsoft Exchange Information Store service is running and that you have the correct permissions to log on. (0x8004011d)

[16:58:52] Errors encountered. Copy process aborted for mailbox 'SuriyaS' ('SURIYAS').


### **Running ExMerge**

Once you have squared away the necessary permissions to run ExMerge against the mailboxes from which you intend to remove a message, you can follow a fairly straightforward procedure to extract the data using the ExMerge program's Archive option. The most effective place to run ExMerge is from the console of the Exchange Server that contains the mailboxes you want to access. However, doing so might require additional rights on the part of your ExMergeUser account, such as being a member of the Remote Desktop Users group on the Exchange Server.

 Some archival systems might immediately archive a copy of a message as soon as they arrive in a user's mailbox. The archival software must be able to extract from the archive messages that should not have been placed there.

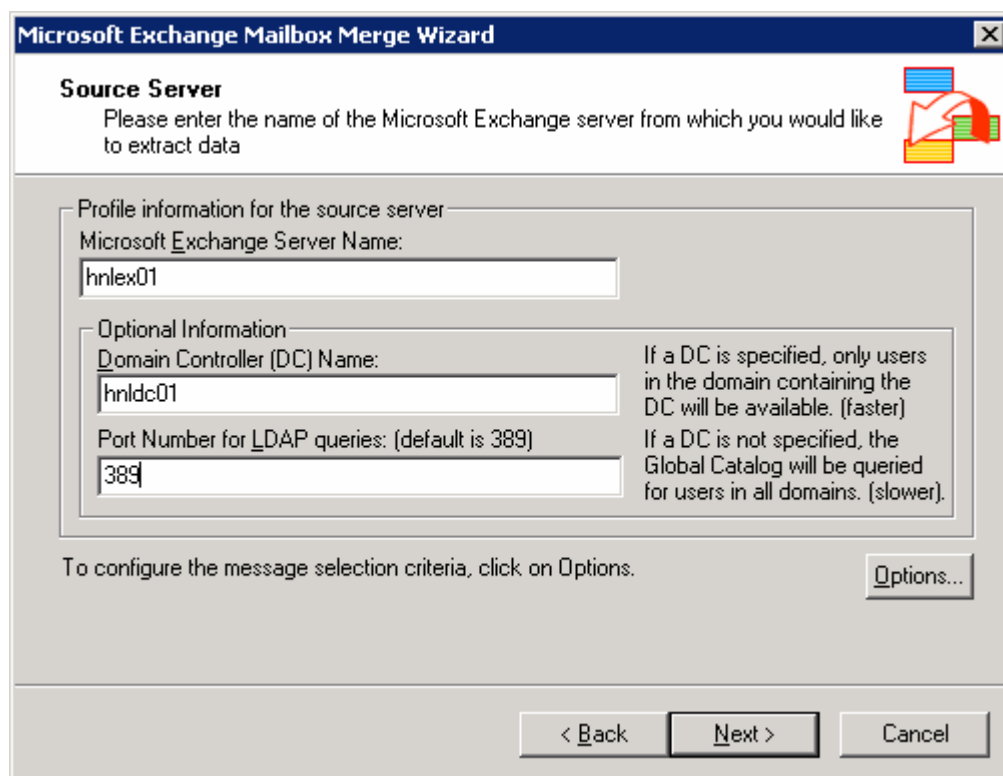
ExMerge allows you to restrict your search for messages using a date/time range, a message subject, or attachment name. If you are attempting to remove a virus or confidential message, you will most definitely want to know the subject or attachment. Using the date/time range is not precise enough to rely upon.

ExMerge will copy or move data from an Exchange Server mailbox to PST files. The PST files will generally take up quite a bit more disk space than the data did when it was in the Exchange Server mailbox store, so allow for plenty of disk space.

 ExMerge creates Outlook 97 to 2002 compatible PST files; it does not create the Office Outlook PST file that provides more than 2GB of storage capacity and is used with Outlook 2003 or later.

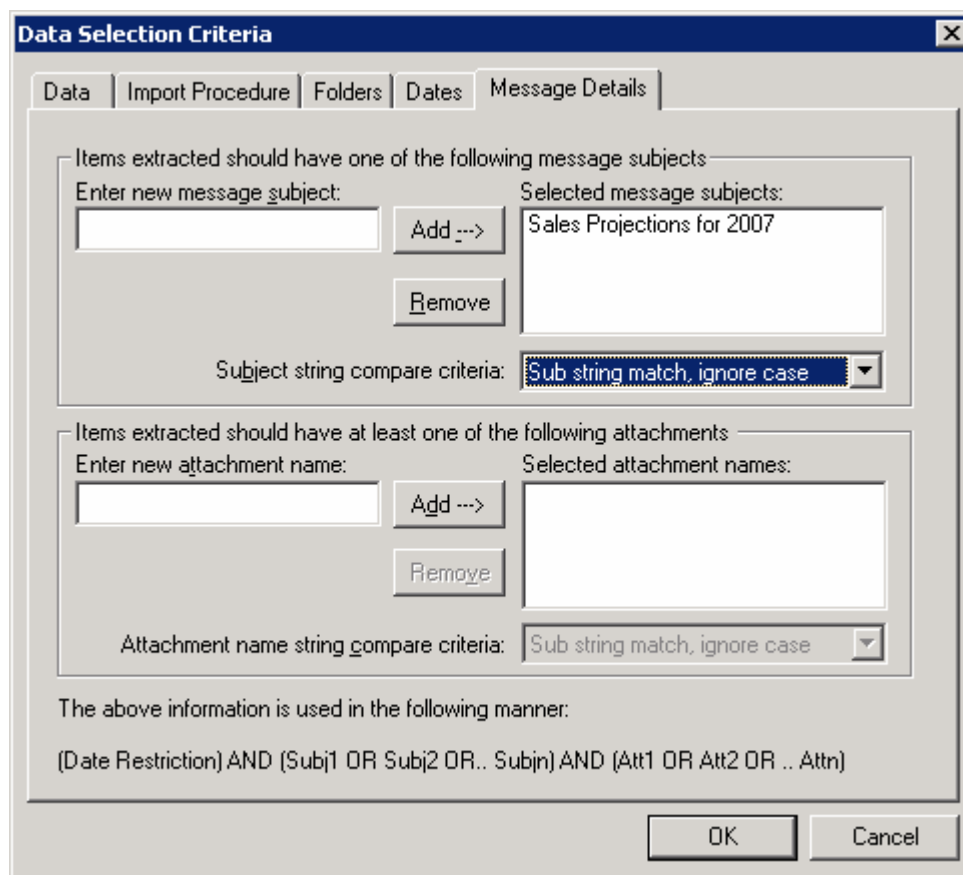
The following example walks through the process of extracting a message that was accidentally sent to many users on a mailbox server. The message contained sensitive information and should not have been sent to the entire list, and its subject line is “Sales Projections for 2007.” This procedure will remove the message from all the mailboxes on the server:

- Run ExMerge, and click Next on the first page of the wizard.
- Choose the Extract or Import (Two Step Procedure) option, and click Next.
- Choose the Step 1: Extract data from an Exchange Server Mailbox, and click Next
- On the Source Server page, provide the name of the Exchange Server from which you plan to export mail and the name of a domain controller. The Lightweight Directory Access Protocol (LDAP) port number is optional. Figure 2.5 shows this dialog box.



**Figure 2.5: Selecting a source Exchange Server and domain controller.**

- Click Options to display the Data Selection Criteria dialog box.
- On the Data tab, you only need to have the User Messages And Folders options.
- On the Import Procedure tab, choose the Archive Data To Target Store option.
- On the Message Details tab (shown in Figure 2.6), enter the message subject in the Enter New Message Subject text box, and click Add. Ensure that the criteria you are going to be extracting by is properly entered. If you don't specify a criterion on the Message Details property page, you will export all the message data from your users' mailboxes—they won't be amused.



**Figure 2.6: Specifying the subjects or attachment names by which you will extract messages.**

- Click OK to close the Data Selection Criteria dialog box and return to the Source Server property page, and click Next.
- On the Database Selection property page, select each of the check boxes corresponding to each database that you want to search through for the message previously specified, then click Next.
- On the Mailbox Selection property page, select the mailboxes that you want to search through. When finished, click Next.
- On the Locale Selection property page, click Next.
- On the Target Directory property page, specify a location for the PST files that will be created. When finished, click Next twice to start the ExMerge process.
- When ExMerge is finished, click Finish.

---

## **Topic 3: Architecture and Deployment Considerations**

### **Q 3.4: What are best practices to follow when deploying servers?**

**A:** Good email security should start the day you take your hardware and software out of the box and start installing. Security should not be an afterthought once a server is in production. Good security and good operational practices go hand-in-hand, and the steps that you take during deployment will contribute to not only better security but also a more stable email server environment.

First and foremost, start with a server build checklist that includes steps that will help make the server more stable and more secure. Once the server operating system (OS) and the email server software is installed, have a checklist of procedures and best practices that you follow to ensure the software is installed properly and securely.

I have a four-layer model that I use for representing the path to stable, secure, and available server platforms. Each of the higher layers depends on the lower layers being built properly. If the lower layers are not built properly, you cannot expect the upper layers to provide reliable or secure services. The four layers are:

- Operational policies and procedures
- Application
- Operating system
- Server platform

### ***Building a Solid Server Platform***

The hardware and the OS are the base for solid operations and good security. Although each server's role may be different and the application software may be different, a consistent installation checklist provides the important first steps. The following checklist should be followed for all server builds. Although there are exceptions to many of these checkpoints, this is a good starting point:

- The server should be placed in a physically secured area with appropriate access and environmental controls (air conditioning and redundant power).
- Server hardware should be updated to have newer versions of the firmware/FlashBIOS, including updates for the system BIOS, disk controllers, host bus adapters (HBAs), tape drives, backplane firmware, and embedded systems management hardware. Newer versions does not always mean the most recent version, so consult with your hardware vendor to confirm the best versions of firmware for your particular environment and application. For Windows Server 2003 (WS2K3), take the default settings for the OS.
- Add only the necessary additional services or Windows components to support the application that will run on the server. Do not install unnecessary components.

- Obtain newer versions of the vendor's server installation CD-ROMs or disks. Confer with the vendor to ensure you are getting the best version for your environment and the application that the server will be supporting.
- Install the OS using the hardware vendor's installation procedures and disks.
- Apply the most recent OS service packs and critical updates for your server platform.
- Consult an automatic update service, such as Microsoft Update, to download the latest critical updates and notify the administrator that they are ready to apply.
- Install the WS2K3 Support Tools.
- Tune the virtual memory settings so that initial page file size is 1.5\*RAM and the maximum file size is 2.0\*RAM.
- Enable disk performance monitor counters (diskperf -y).
- Optimize the WS2K3 memory settings based on your intended application. For Exchange Server 2003, see Microsoft Knowledge Base article 815372 "How to optimize memory usage in Exchange Server 2003."
- Apply the appropriate baseline security template to the server based on your environment. Once the template is applied, test the server and its functions thoroughly before proceeding to ensure that the template has not broken a critical function that you require.
- Configure the Windows event log sizes, such as the following:
  - System log = 49,152KB
  - Security log = 49,152KB
  - Application = 196,608KB
- Install the application, such as Exchange Server 2003, and apply the most recent service pack and critical updates.
- Install antivirus software that supports the antivirus APIs provided by your mail server software, such as the Exchange antivirus application programming interface (AVAPI). Configure the antivirus software to update signatures between 4 and 12 times per day; doing so ensures that you check for updated virus signatures at least once every 6 hours and possibly as frequently as once every other hour.
- Disable any Windows services that you have determined are unnecessary for your organization.
- If your Exchange OWA is not protected by a reverse proxy server that can filter URLs and allow only specific virtual directories to access the Exchange Server, consider implementing the Urlscan utility. See Microsoft Knowledge Base article 823175 "Fine-tuning and known issues when you use the Urlscan utility in an Exchange 2003 environment" for more information. Keep in mind that Urlscan may disable a user's ability to open messages with special characters in the subject line.



## Checking Your Work

After you have your server initially installed and configured, you can proceed to check your work and make sure you have not missed anything obvious or potentially harmful. For Windows and Exchange, Microsoft provides two simple-to-use tools for just this purpose—the Microsoft Baseline Security Analyzer 2.0 (MSBA) and the Exchange Server Best Practices Analyzer Tool (ExBPA).

The MSBA is a more generic tool that scans to determine whether critical updates are available for programs such as Windows, Internet Explorer (IE), Exchange, and Office as well as scans for vulnerabilities to well-known problems. When you run the MSBA, it connects to Microsoft and downloads the latest XML file it will use to scan for vulnerabilities and patch versions. The scan report includes information about what was scanned and how to correct potential problems. Figure 3.20 shows an MSBA report that was run against a WS2K3 member server supporting Exchange Server 2003.

The screenshot displays the Microsoft Baseline Security Analyzer 2.0 interface. The main window shows a security report for a scan performed on 3/15/2006 at 1:48 PM. The scan was conducted using MBSA version 2.0.5029.2 against the Microsoft Update catalog. The overall security assessment is 'Severe Risk' due to one or more critical checks failing.

The report is sorted by 'Score (worst first)' and is divided into several sections:

- Windows Update Scan Results:** A table with columns for Score, Issue, and Result. The result shows 'No security updates are missing' with a score of 100.
- Windows Scan Results:** A section containing 'Administrative Vulnerabilities'.
- Administrative Vulnerabilities:** A table with columns for Score, Issue, and Result. It lists several issues:
 

Score	Issue	Result
0	Automatic Updates	Updates are automatically downloaded, but not automatically installed on this computer.
0	Password Expiration	Some user accounts (2 of 6) have non-expiring passwords.
0	Incomplete Updates	No incomplete software update installations were found.
0	Windows Firewall	Windows Firewall is not installed or configured properly, or is not available on this version of Windows.
100	Local Account Password Test	Some user accounts (1 of 6) have blank or simple passwords, or could not be analyzed.
100	File System	All hard drives (1) are using the NTFS file system.
100	Autologon	Autologon is not configured on this computer.

Figure 3.20: Baseline Security Analyzer report.

The second tool that is helpful in isolating a poorly configured server is ExBPA. This tool is more targeted towards analyzing Exchange Servers but will look at some configuration items related to domain controller/Global Catalog (GC) configuration as it pertains to Exchange Server. The ExBPA uses WMI to collect information about one or more of your Exchange Servers and analyzes this against a database of best practices and configurations that is a result of the experience of the Exchange team as well as many consultants and Exchange experts. Regardless of how well you think your Exchange organization is configured, the ExBPA will probably find something you have not thought of.

When you first launch the ExBPA, it will want to download any updates to its XML database of best practices and will then need to connect to Active Directory (AD) to read your Exchange organizational structure. Figure 3.21 shows the initial screen of the ExBPA in which you are asked which Exchange Servers are to be analyzed and the type of scan to perform.

## Start a New Best Practices scan

Enter an identifying label for this scan:

---

Specify the scope for this scan:

- Volcano Surfboards
  - First Administrative Group
    - HNLEX01
    - KILAUEA

Summary:

Scope is set to 2 server(s), 1 administration group(s), and the organization.

---

Select the type of scan to perform:

---

Select the speed of the network to adjust estimated time value:

This scan will take approximately 8 minutes

---

**Figure 3.21: Starting an ExBPA scan.**

Once you click the Start Scanning option on the main page, ExBPA will start scanning AD and the selected servers. This scan may take anywhere from a few minutes to a few hours depending on the number of Exchange Servers you have selected to scan and the network bandwidth between the scanning machine and the selected servers.

When the report is completed, you can sort the report based on a number of criteria. The default is the critical issues list, but I usually switch to the Full Issues List report so that I can list everything that the ExBPA found. The Full Issues List is shown in Figure 3.22. Each issue can be expanded to see information about what the recommended setting is and if you would like the ExBPA to skip that particular setting the next time it scans that server.

ExBPA reports might contain a considerable amount of information about your organization's internal infrastructure. Treat these reports as confidential and protect them accordingly.

A word of caution about these reports: the best practices recommendations are generic and don't take into consideration your organization's message flow rules or your business practices. For example, the report in Figure 3.22 recommends "Consider setting 'TarpitTime.'" However, in this organization, all inbound mail is routed through a managed provider and the organization's firewall only accepts inbound SMTP from the managed provider's SMTP servers. Thus, configuring an SMTP tarpit will not have any effect on the organization's messaging services.

## View Best Practices Report

March 2006 BPA Scan

Select a report:

[Find](#) [Export report](#) [Print report](#)

**Full Issues List** Arrange by:

Total number: 33 items found

**Organization Volcano Surfboards**

- ⚠ Policy change required Organization: Volcano Surfboards
- ⚠ RUS did not process all changes Organization: Volcano Surfboards
- ⚠ RUS did not process all changes Organization: Volcano Surfboards

**Server HNLEX01**

- ❌ VMware detected Server: HNLEX01
- ⚠ SMTP performance warning Server: HNLEX01
- ⚠ Recovery storage group enabled Server: HNLEX01
- ⓘ Application log size Server: HNLEX01  
 As a best practice, the size of the 'Application' log on server HNLEX01.volcanosurfboards.com should be increased. The current size is 16MB. For servers running Microsoft Exchange, a size of 40MB or more is recommended.
  - ➔ Do not show me this item again for this instance only.
  - ➔ Do not show me this item again for all instances.
- ⓘ Consider setting 'TarpitTime' Server: HNLEX01
- ⓘ Enable IME automatic updates Server: HNLEX01

Figure 3.22: An example ExBPA Full Issues List report.

An overzealous administrator may also implement these recommendations without fully considering the ramifications. A common recommendation is to disable anonymous authentication SMTP on Exchange back-end servers. If you don't read the recommendation carefully, you might disable SMTP authentication on your front-end servers and begin rejecting mail from the Internet. Exercise caution, read the recommendations carefully, and consider how they will affect your organization.

### **Infrastructure Protection**

For servers that are connected to an outside network (such as the Internet or to a business partner), there are additional steps you can perform to provide another layer of protection for your servers:

- Use a managed provider on the Internet or SMTP message hygiene system in the DMZ to provide initial inspection of message content.
- Protect published Web resources such as Outlook Web Access (OWA) or mobile device access (ActiveSync) by placing a reverse proxy server between the internal servers and the Internet. Ensure that the reverse proxy server can provide URL inspection.
- On Internet-facing servers, use SSL certificates that have been issued by a trusted certifying authority.
- Configure OWA to use Forms Based Authentication (FBA) so that logon sessions are timed out after some period of client inactivity.
- Implement IPSec between domain controllers/GC servers and Exchange Servers.

### **Bad Practices**

To the best practices list, there is a corresponding list of things that you should not do or that may affect the stability or security of your messaging system:

- Installing a mail client on the console of the mail server
- Surfing the Web, checking email, or using desktop applications from the server console
- Installing unnecessary Windows components or additional applications on dedicated mail servers
- Installing evaluation software on production servers
- Combining server roles such as domain controller, Exchange Server, and SQL on the same physical machine
- Applying service packs and critical fixes immediately after they are released and without testing them

---

## **Topic 4: Protecting and Controlling Sensitive Information in Email**

### **Q 4.4: What options are available for antivirus and spam protection?**

**A:** One of the top issues that IT management now reports is the challenge of reducing spam and preventing the spread of hostile content such as viruses and worms. IT management is faced with the overwhelming cry of “Stop this spam!” from the user community and the almost equally loud cry of “Keep viruses and worms off the mail system” from the information security staff.

The number of solutions on the market is almost staggering and each claims to be the ultimate solution to the spam or virus problem. In the past few years, a convergence and consolidation of technologies has happened in the field of message hygiene. Now, most products offer a cradle-to-grave solution for fighting spam and viruses rather than requiring separate components, appliances, services, or software packages.

#### ***Managed Providers***

In the past few years, the trend towards managed providers in the fight against unwanted email content has taken a positive turn. Managed providers handle all your incoming mail by having your DNS MX records for your Internet domain pointed to their Simple Mail Transfer Protocol (SMTP) servers. These servers perform message hygiene/content inspection on messages intended for your mail system, then pass the mail on to your mail system.

IT managers who not long ago would have considered it an unacceptable practice to have their email relayed through third party are now embracing the concept. There are clearly several advantages to using a managed provider:

- Managed providers are usually staffed 24 × 7.
- Managed providers offer almost immediate response to day-zero threats.
- Managed providers have better scalability and, often, multiple locations to provide fault tolerance.
- An organization’s mail server can be configured to accept only inbound mail from the provider’s SMTP server, thus reducing the potential for a Denial of Service (DoS) attack against that organization’s SMTP servers.
- Spam and viruses are filtered by the managed provider and do not ever enter the organization’s DMZ or mail servers. Thus, bandwidth, disk space, and system resources are not consumed by unwanted email.
- In the event of Internet connection or mail server failures within an organization’s network, the managed provider can queue mail until service is restored.

One weakness of some of these providers is that they may not have a listing of your valid email addresses and thus will not reject mail for your invalid recipients. Even for a small organization with spam problems, this shortcoming can generate a lot of inbound mail traffic as spammers use dictionary spam methods to send mail to random names. When choosing a managed provider, confirm that they can synchronize with your directory or that you can upload a list of valid email addresses to their system.

### ***Protection in the DMZ***

A fairly common method of protecting SMTP mail systems is to provide an SMTP relay server in the organization's DMZ or perimeter network. All inbound mail is delivered to the SMTP relay. The SMTP relay in the DMZ performs the message hygiene functions.

In the past, this SMTP relay may have been something as simple as a UNIX or Windows server that performed only message forwarding; the system's hardware may have been an Intel or RISC machine. These systems are potentially far more complex than they were just a few years ago. Although Windows- and UNIX-based servers may still run the software, it now consists of complex message hygiene functions that perform virus detection and spam filtering. The message hygiene system may run on a regular server, but a more common solution is for vendors to package the software on vendor-provided hardware that fits in a 1U or 2U rack space. These software/hardware appliances still run a hardened version of Windows or UNIX, but they hide the additional complexities of these machines from the customer.

The advantage of providing an initial point of message hygiene in the DMZ network is that all inbound mail is inspected before it ever makes it to the mailbox server. The message hygiene system also protects the internal mail servers by ensuring that no inbound SMTP connection from the Internet arrives directly on the mailbox servers.

### ***Protection on the Mail Server***

Message hygiene systems that run directly on the mail server have been around for several years. Some of these systems are not truly integrated with the mail system and simply perform SMTP content inspection on inbound messages before passing the messages on to the mail server's SMTP service. Other systems integrate with the mail system's SMTP service or message storage system.

The Exchange 2003 antivirus API (AVAPI) enables message hygiene systems designed for Exchange to inspect messages as they arrive via the Windows SMTP service or once the message is moved into the information store. AVAPI software will also prevent a user from opening a message until it has been inspected by the message hygiene software.

Although installing message hygiene software directly on the mail server is an important part of a good security practice, it should be considered a second-tier defense. Multiple layers of protection are important. The first line of defense should be an inspection system provided by a managed provider or in the organization's DMZ.

---

## **Topic 5: Firewall Strategies and Best Practices**

### **Q 5.4: What are best practices for configuring a firewall to support better email security?**

**A:** Your organization's firewall is the first line of defense against intruders, Denial of Service (DoS) attacks, and unwanted messaging content. Depending on how you look at things, the firewall is also the Internet's first line of defense against your network as well.

There are several actions you can take when configuring your Internet-facing and internal firewalls (if applicable) to provide better messaging security:

- Configure a DMZ or perimeter network to support Simple Mail Transfer Protocol (SMTP) content-inspection systems or HTTP reverse proxy servers.
- Open only the necessary ports between services in the DMZ and the internal mail system resources. For example, don't open inbound port 25 to all internal SMTP servers if the only servers that should be accepting inbound SMTP mail are the Exchange front-end/bridgehead servers.
- Block outbound SMTP message traffic so that only authorized hosts can send SMTP mail to external hosts. Doing so will prevent a client that may become infected with a worm that has its own SMTP engine from sending mail to the Internet. Having your ISP or a business partner call you to inform you that one of your computers is sending out a virus is pretty embarrassing. Most users should never have direct outbound SMTP access. If this is required, have them configured to use your content inspection systems for outbound SMTP relay.
- If possible, provide SMTP application filters that inspect inbound SMTP packets for potential DoS attacks, directory harvesting, or other hostile SMTP traffic.
- If you use a managed provider, configure inbound SMTP to be accepted only from the IP addresses of your managed provider's servers. Your managed provider can give you this list as well as notices of changes to the list.
- Confirm that the firewall's external IP addresses have DNS PTR records created with the owner of the reverse lookup zones. These IP addresses are often evaluated by remote SMTP servers and potentially rejected if there is not a valid PTR record.



---

## **Topic 6: Protecting and Controlling Sensitive Information in Email**

### **Q 6.4: How does Enterprise Rights Management work?**

**A:** On first glance, the big picture for Enterprise Rights Management (ERM) is deceptively simple. ERM helps protect sensitive content and helps an organization enforce information security policies. When you start digging into the first layer of technical details, though, ERM systems may appear to be overly complex and difficult to manage; furthermore, the system may appear to be difficult enough to use that typical corporate users would refuse to use the solution.

At least that was my initial impression, but upon learning more about ERM solutions from both the end-user perspective and the configuration, the complexity was merely a lack of basic understanding of the concepts and the flow of rights and information. And much to my relief, the complexity of ERM solutions is hidden from the end user. Using an ERM solution is not much more difficult than saving a document.

For an IT administrator or a messaging administrator intent on providing better information protection and information security policy enforcement, a basic understanding of the concepts and the inner workings is helpful when deploying or supporting ERM solutions. Keep in mind that there are several vendors with ERM solutions on the market. Each of these has strengths and weaknesses and the components are set up differently. For my example of how an ERM solution works, I'm using the Microsoft Windows Rights Management Server for Windows Server 2003 (WS2K3), the Windows Rights Management client, and Office 2003 Professional applications.

### ***Publishing Licenses and Usage Licenses***

Different ERM solutions handle the creation of Publishing Licenses and Usage Licenses differently. The Publishing License is created when the content publisher sends the encrypted content keys and rights information to the Rights Management Server (RMS) Licensing server and the RMS system signs that information. The Usage License is created when a content consumer sends the Publishing License to the RMS Licensing server along with the consumer's Rights Account Certificate; the RMS Licensing server creates the Usage License that will allow the user to open this content.

Both of these pieces are key components to any ERM solution; however, the storage of these components, usage, and issuance of the Usage Licenses may differ from system to system. Some vendors rely exclusively on a policy server to store the rights information, while others (such as Microsoft) embedded the Publishing License into the protected content; once the user has received a Usage License, that license is also embedded in the content. This setup allows the content consumer to use the content offline and eliminates a potential point of failure.



## RMS

The core of most ERM solutions is the RMS. In a single RMS-system organization, the RMS handles two primary functions. The first function is user certification; when a user uses the RMS client software for the first time to either consume protected content or create protected content, the RMS' certification function is responsible for issuing new user certificates (also known as an account certificate).

The second responsibility of an RMS is through the licensing services. Licensing services are responsible for issuing *publishing licenses* when protected content is created and issuing *use licenses* when protected content is consumed. In a larger environment, additional licensing servers may be configured for fault tolerance or to provide closer licensing services to users across slow WAN links. Figure 6.3 shows an organization with an RMS root server and two additional licensing servers.

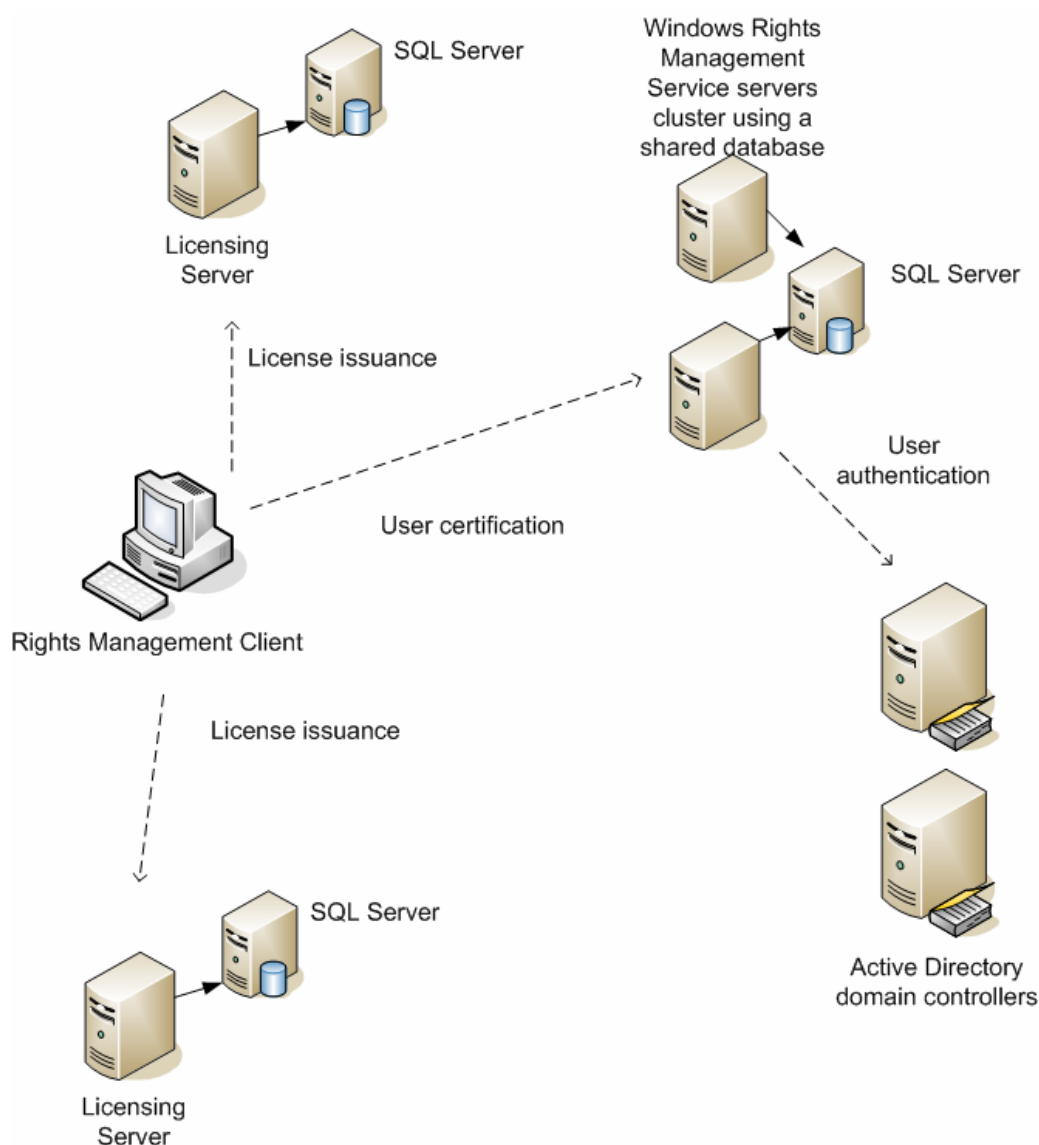


Figure 6.3: Components of a Windows RMS setup.

In a smaller organization, a single RMS would handle both user certification and licensing. In larger organizations, there could be many licensing servers and many RMSs in an RMS cluster. An RMS cluster is multiple RMSs that share a single database and can be scaled by adding additional RMSs that all share a common SQL Server database.

RMSs use Active Directory (AD) to authenticate all users of the RMS system when issuing certificates or enrolling users. AD stores the service connection point (SCP) which directs users to their RMS. AD is also used to resolve permissions that are assigned to distribution lists; this is necessary when content is protected and assigned for use by distribution groups.

### **Publisher Setup**

In an RMS system, the creator of protected content (or at least the person that applies the rights to content) is known as the “publisher.” Before the publisher can create protected content, a number of prerequisites have to be met—RMS-enabled applications and the client software must be installed and the user and the machine must be licensed. The process of setting up the client is called bootstrapping.

### **RMS-Enabled Applications**

First and foremost, before protected content can be created or consumed, the user must have RMS-enabled applications. Office 2003 Professional applications such as Microsoft Word, Excel, PowerPoint, and Outlook are already enabled and ready for use. If Web-based content needs to be protected, the HTML content can be created using Word 2003 and viewed with Internet Explorer (IE) and the Rights Management Add-on for IE.

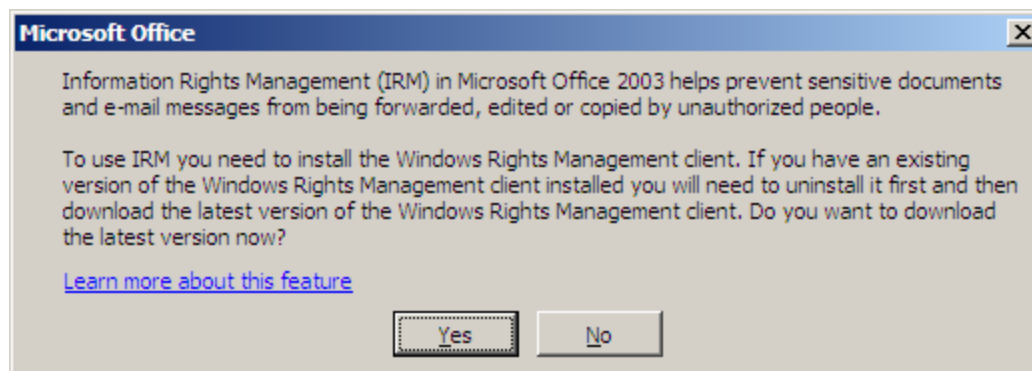


If you are supporting applications other than the Office 2003 Professional suite, such as Office XP or earlier versions of the Office Suite or Adobe Acrobat, vendors such as Liquid Machines and GigaTrust have additional tools to protect content consumed by these applications.

If you have specialized applications that are not covered by the principal vendors, the Rights Management Software Development Kit (SDK) can be used to extend your custom applications to enable RMS features and protection.

### **Rights Management Client Software**

All Windows RMS applications must interact with a common piece of client software. Although rights management functions are built-in to the Windows Vista operating system (OS), earlier versions of Windows must use the Windows Rights Management Service (WRMS) client; the current version is WRMS Service Pack (SP1), which can be installed directly (you don't need to install the original version first). If a user tries to set RMS permissions on a file without the RMS client, the user will receive a notice informing them that the WRMS client needs to be installed (see Figure 6.4).

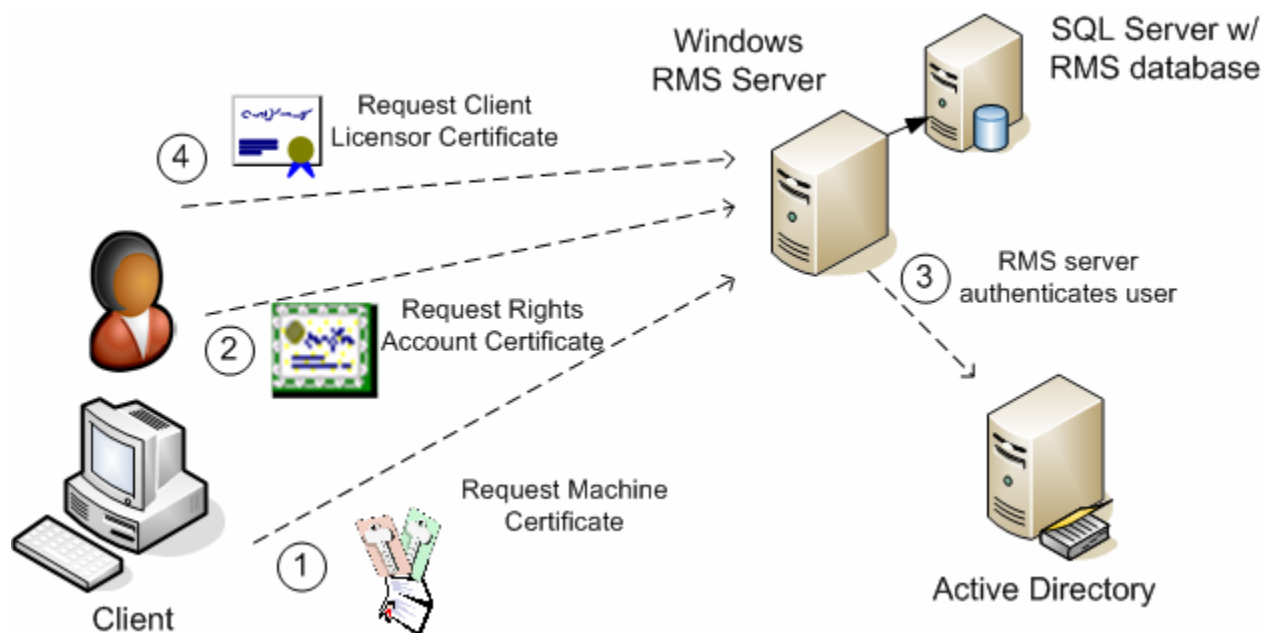


**Figure 6.4:** The user is warned if they don't have RMS client software.

The WRMS client can be installed on Windows 2000 (Win2K) SP4 and later, Windows XP SP1 and later, and WS2K3. Prior to any RMS-enabled application being used, the client software must be installed; there are no configuration options and there are no user-executable components to this software. The software does not need to be activated or used at installation, so it can be installed when the OS and applications are installed or as part of an imaging process.

### Using RMS for the First Time

Before users can create and consume content, they must be trusted by the RMS system. There are multiple parts to this process, including having a Machine Certificate issued to the computer as well as certificates issued to the user that allow the user to publish both online and offline. The RMS client software has to locate an RMS using the SCP information in AD. The steps to user provisioning are conceptually shown in Figure 6.5.



**Figure 6.5:** RMS user provisioning process.

In the first step, the RMS client software generates a public/private key pair, connects to the RMS, and requests that the public key be signed so that a Machine Certificate can be created. One Machine Certificate is issued per user per machine. The Machine Certificate is used to authenticate the computer in the RMS environment and to unlock users' private keys upon request. This process is completely transparent to the user (and the administrator) and does not require that the user have administrative permissions on the computer.

The second step of the user provisioning process requires that the user become trusted by the RMS system by obtaining a Rights Management Account Certificate. This process requires that the user provide domain credentials; the RMS (step 3 in Figure 6.5) authenticates the user against AD.

Once authenticated, the user's public/private key pair is generated, the Rights Management Account Certificate (containing the user's public key) is created, the Rights Management Account Certificate and the private key are stored in the SQL database, and the user's private key is encrypted with the public key found in the Machine Certificate. This extra step ensures that the user's private key remains protected and ensures that the only way the user can create or consume RMS-protected content is from a trusted workstation. The user is now trusted by the RMS system and can create or consume content.

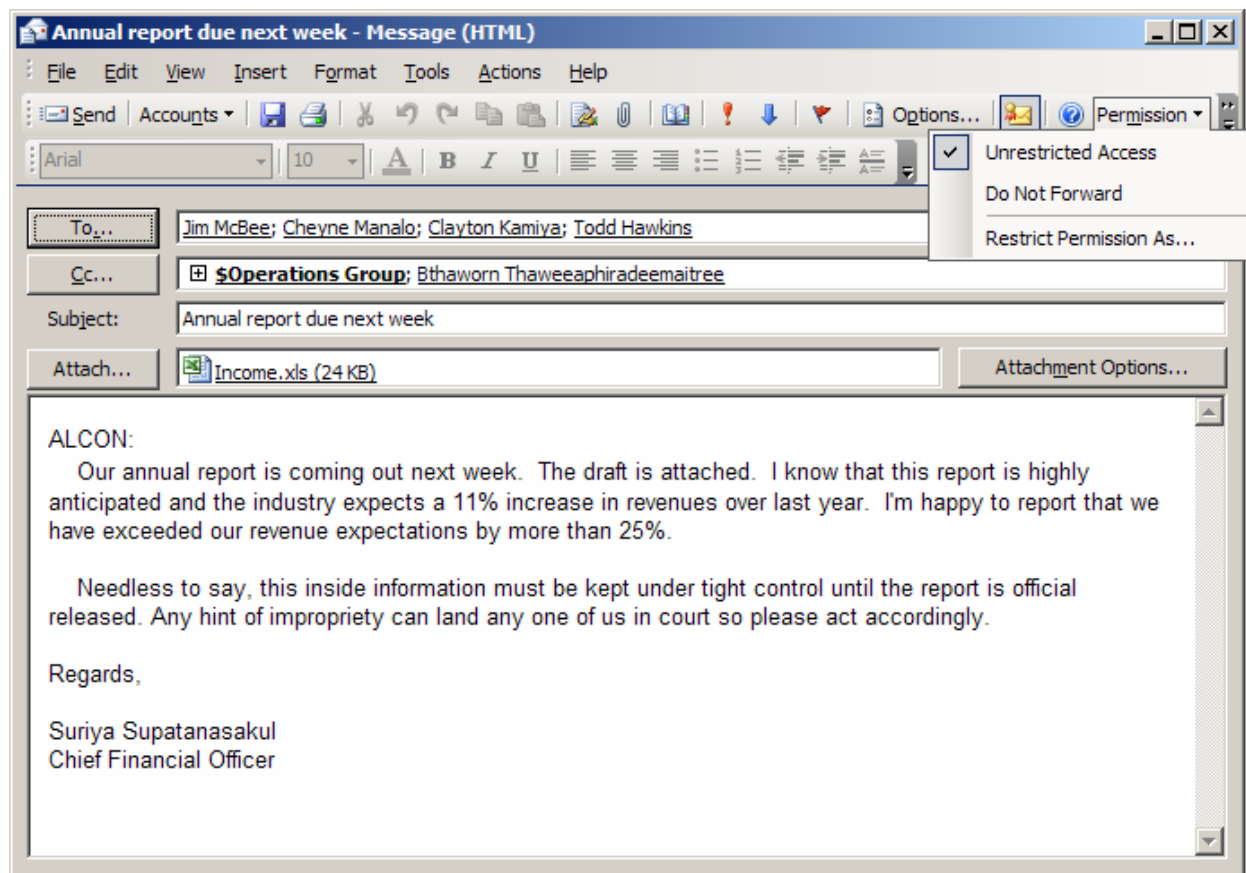


Except for being challenged for credentials, the Rights Management Account Certificate creation process is completely transparent to the end user.

Finally, in step 4, the user is issued a Client Licensor Certificate (CLC). The CLC is used to allow the user to create protected content without having to connect to the RMS. The CLC is signed by the RMS' public key and contains the CLC public key, the CLC private key (encrypted by the user's Rights Management Account Certificate public key), and a copy of the RMS' Licensor Certificate. The CLC will enable the user to publish protected content on behalf of the RMS.

### ***Creating Protected Content***

Now that both the machine and the user are trusted by the RMS system, the user can create protected content. Microsoft Office 2003 Professional applications all have a File, Permissions option on their menus; the Permissions option can also be placed on a toolbar if it is frequently used. Figure 6.6 shows Office 2003 and the Permissions options; to protect this sensitive message with rights management, the user must choose the Restrict Permissions As option. As the contents of this message reveal, this message is something that the originator clearly does not want disclosed either intentionally or accidentally. Protecting the message contents via ERM will impose the necessary restrictions on the message.

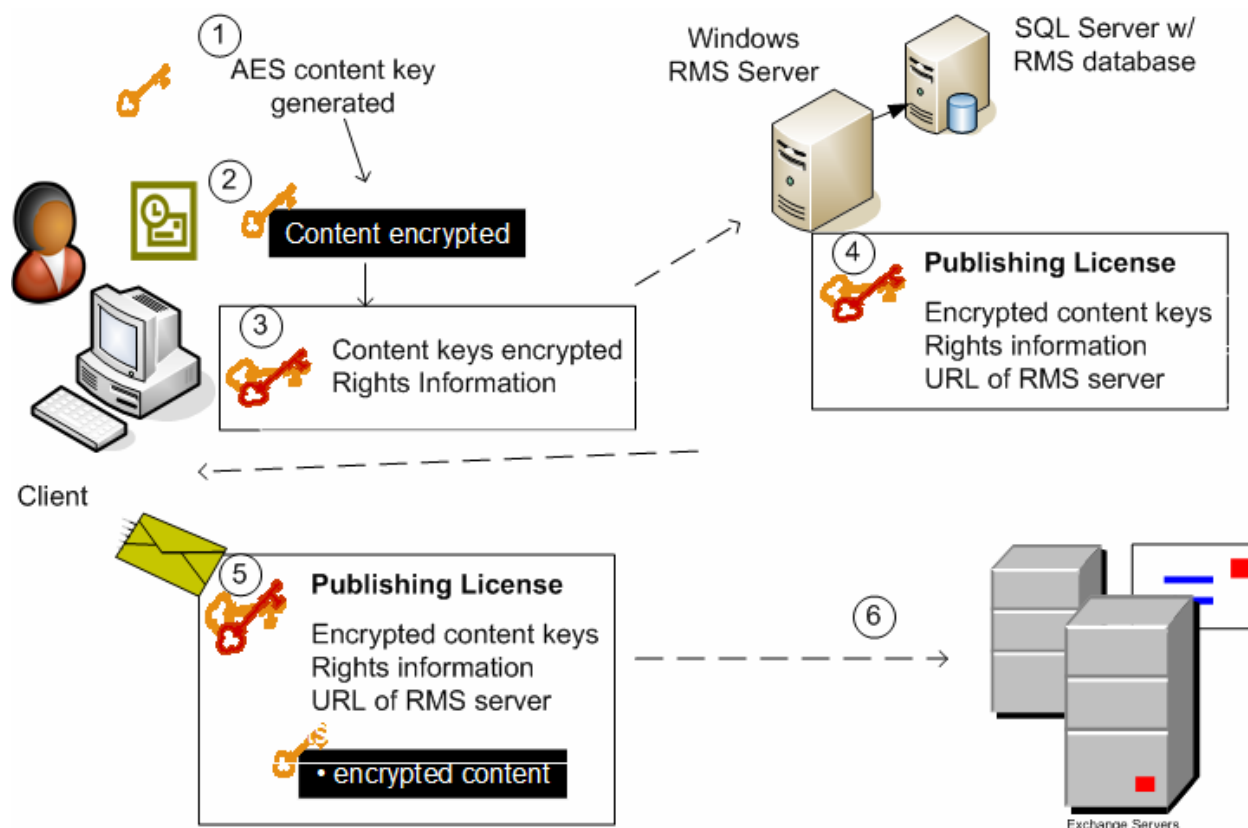


**Figure 6.6: Protecting a sensitive message.**

When the user chooses the Restrict Permissions As option from the Permissions menu, the application (in this case, Outlook 2003) will go through the necessary steps to protect the content. These steps are conceptually illustrated in Figure 6.7. Although the process in Figure 6.7 seems fairly complex, all the user had to do was to choose an option from the menu.

 All RMS client to RMS server communication takes place over HTTP or HTTPS.


In the first step, the RMS client generates an advanced encryption standard (AES) content key (this is a symmetric key); this key is used to encrypt the contents of the data to be protected. Next, the RMS client software encrypts the AES key twice; once with the user's public RMS key and once with the RMS' public key. The encrypted content keys and the rights information are then sent to an RMS; in a midsized or large RMS environment, this may not be the same RMS that provisioned the user but instead a dedicated RMS Licensing server.




**Figure 6.7: Creating RMS-protected content.**

When the encrypted content keys and rights information is received by the RMS Licensing server, the RMS Licensing server verifies that the user is trusted by the RMS system, creates the publishing license, and signs it with the RMS' private key. The Publishing License is then retrieved by the client.

The RMS-enabled client software appends the Publishing License to the encrypted content. In the case of a document or spreadsheet, the file is then saved with the Publishing License and the encrypted content. The file is now protected regardless of whether it is on the local file system, optical media, or a USB storage device. In the example in Figure 6.7, the content is an email message; the message is sent on to the email server and ultimately to the intended recipients.

 With the Windows RMS solution, security policy information is distributed with the protected content. Only RMS-enabled applications can open the protected content. The RMS-enabled application enforces the rights restrictions.

All the content publisher had to do is set the permissions option, and the appearance of the message is altered slightly (see Figure 6.8); in the area above the To line of the message, the rights information is displayed (recipients cannot forward, print, or copy the content).




**Figure 6.8:** The message display changes once the message is protected.

The procedure illustrated in Figure 6.7 uses the Rights Account Certificate (RAC) and requires connectivity to an RMS Licensing server in order to create protected content. If the publisher is working offline and has a Client Licensor Certificate (CLC), they can create protected content on behalf of the RMS.

### **Consuming Protected Content**

The final part of the puzzle is when the recipient receives the protected content; the content must be decrypted and the rights information examined so that the RMS-enabled application can present the recipient with the correct rights to the content. When the RMS-enabled application opens the message, the only information that the application can see is the Publishing License containing the encrypted content keys, the rights information, and the URL of the RMS Licensing server.

 Based on the Publishing License, the only two private keys that can open the encrypted content key and thus decrypt the protected content are the user that created the content and the RMS' private key.

The RMS client software takes the Publishing License and its RAC and sends a request for a Use License to the RMS Licensing server via HTTP or HTTPS. The RMS Licensing server validates the requestor's RAC, inspects the Publishing License for the rights list, and validates the requesting user in AD.

 An RMS must be available on the network in order for a Use License to be issued.

The RMS then uses its private key to decrypt the AES content encryption key and then encrypts the AES content encryption key with the recipient's RAC public key. The Use License is created that contains the AES encrypted key and is returned to the recipient. The RMS-enabled application can now open the content and will enforce the published rights. In the case of a Word document or Excel spreadsheet, the Use License is embedded into that recipient's copy of the document. For Outlook 2003, though, the Use License is stored in the user's local user profile directory.



## Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.