CIS Exchange Server 2003 Benchmark
Version 1.0

This Page Intentionally Left Blank

Warnings

Do not attempt to implement any of the settings in this guide without first testing in a non-operational environment.

This document is only a guide containing recommended security settings. It is not meant to replace well-structured policy or sound judgment. Furthermore, this guide does not address site-specific configuration issues. Care must be taken when implementing this guide to address local operational and policy concerns.

The security configuration described in this document has been tested on a Windows Server 2003 system. Extra care should be taken when applying the configuration in other environments.

SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE EXPRESSLY DISCLAIMED. IN NO EVENT SHALL THE CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

TERMS OF USE AGREEMENT

Background.

The Center for Internet Security ("**CIS**") provides benchmarks, scoring tools, software, data, information, suggestions, ideas, and other services and materials from the CIS website or elsewhere ("**Products**") as a public service to Internet users worldwide. Recommendations contained in the Products ("**Recommendations**") result from a consensus-building process that involves many security experts and are generally generic in nature. The Recommendations are intended to provide helpful information to organizations attempting to evaluate or improve the security of their networks, systems, and devices. Proper use of the Recommendations requires careful analysis and adaptation to specific user requirements. The Recommendations are not in any way intended to be a "quick fix" for anyone's information security needs.

No Representations, Warranties, or Covenants.

CIS makes no representations, warranties, or covenants whatsoever as to (i) the positive or negative effect of the Products or the Recommendations on the operation or the security of any particular network, computer system, network device, software, hardware, or any component of any of the foregoing or (ii) the accuracy, reliability, timeliness, or completeness of the Products or the Recommendations. CIS is providing the Products and the Recommendations "as is" and "as available" without representations, warranties, or covenants of any kind.

User Agreements.

By using the Products and/or the Recommendations, I and/or my organization ("**We**") agree and acknowledge that:

1. No network, system, device, hardware, software, or component can be made fully secure;

2. We are using the Products and the Recommendations solely at our own risk;

3. We are not compensating CIS to assume any liabilities associated with our use of the Products or the Recommendations, even risks that result from CIS's negligence or failure to perform;

4. We have the sole responsibility to evaluate the risks and benefits of the Products and Recommendations to us and to adapt the Products and the Recommendations to our particular circumstances and requirements;

5. Neither CIS, nor any CIS Party (defined below) has any responsibility to make any corrections, updates, upgrades, or bug fixes; or to notify us of the need for any such corrections, updates, upgrades, or bug fixes; and

6. Neither CIS nor any CIS Party has or will have any liability to us whatsoever (whether based in contract, tort, strict liability or otherwise) for any direct, indirect, incidental, consequential, or special damages (including without limitation loss of profits, loss of sales, loss of or damage to reputation, loss of customers, loss of software, data, information or emails, loss of privacy, loss of use of any computer or other equipment, business interruption, wasted management or other staff resources or claims of any kind against us from third parties) arising out of or in any way connected with our use of or our inability to use any of the Products or Recommendations (even if CIS has been advised of the possibility of such damages), including without limitation any liability associated with infringement of intellectual

property, defects, bugs, errors, omissions, viruses, worms, backdoors, Trojan horses or other harmful items.

Grant of Limited Rights.

CIS hereby grants each user the following rights, but only so long as the user complies with all of the terms of these Agreed Terms of Use:

1. Except to the extent that we may have received additional authorization pursuant to a written agreement with CIS, each user may download, install and use each of the Products on a single computer;

2. Each user may print one or more copies of any Product or any component of a Product that is in a .txt, .pdf, .doc, .mcw, or .rtf format, provided that all such copies are printed in full and are kept intact, including without limitation the text of this Agreed Terms of Use in its entirety.

Retention of Intellectual Property Rights; Limitations on Distribution.
The Products are protected by copyright and other intellectual property laws and by international treaties. We acknowledge and agree that we are not acquiring title to any intellectual property rights in the Products and that full title and all ownership rights to the Products will remain the exclusive property of CIS or CIS Parties. CIS reserves all rights not expressly granted to users in the preceding section entitled "Grant of limited rights."

Subject to the paragraph entitled "Special Rules" (which includes a waiver, granted to some classes of CIS Members, of certain limitations in this paragraph), and except as we may have otherwise agreed in a written agreement with CIS, we agree that we will not (i) decompile, disassemble, reverse engineer, or otherwise attempt to derive the source code for any software Product that is not already in the form of source code; (ii) distribute, redistribute, encumber, sell, rent, lease, lend, sublicense, or otherwise transfer or exploit rights to any Product or any component of a Product; (iii) post any Product or any component of a Product on any website, bulletin board, ftp server, newsgroup, or other similar mechanism or device, without regard to whether such mechanism or device is internal or external, (iv) remove or alter trademark, logo, copyright or other proprietary notices, legends, symbols or labels in any Product or any component of a Product; (v) remove these Agreed Terms of Use from, or alter these Agreed Terms of Use as they appear in, any Product or any component of a Product; (vi) use any Product or any component of a Product with any derivative works based directly on a Product or any component of a Product; (vii) use any Product or any component of a Product with other products or applications that are directly and specifically dependent on such Product or any component for any part of their functionality, or (viii) represent or claim a particular level of compliance with a CIS Benchmark, scoring tool or other Product. We will not facilitate or otherwise aid other individuals or entities in any of the activities listed in this paragraph.

We hereby agree to indemnify, defend, and hold CIS and all of its officers, directors, members, contributors, employees, authors, developers, agents, affiliates, licensors, information and service providers, software suppliers, hardware suppliers, and all other persons who aided CIS in the creation, development, or maintenance of the Products or Recommendations ("**CIS Parties**") harmless from and against any and all liability, losses, costs, and expenses (including attorneys' fees and court costs) incurred by CIS or any CIS Party in connection with any claim arising out of any violation by us of the preceding paragraph, including without limitation CIS's right, at our expense, to assume the exclusive defense and control of any matter subject to this indemnification, and in such case, we agree to cooperate with CIS in its defense of such claim. We further agree that all CIS Parties are third-party beneficiaries of our undertakings in these Agreed Terms of Use.

Special Rules.

The distribution of the NSA Security Recommendations is subject to the terms of the NSA Legal Notice and the terms contained in the NSA Security Recommendations themselves (http://nsa2.www.conxion.com/cisco/notice.htm).

CIS has created and will from time to time create, special rules for its members and for other persons and organizations with which CIS has a written contractual relationship. Those special rules will override and supersede these Agreed Terms of Use with respect to the users who are covered by the special rules.

CIS hereby grants each CIS Security Consulting or Software Vendor Member and each CIS Organizational User Member, but only so long as such Member remains in good standing with CIS and complies with all of the terms of these Agreed Terms of Use, the right to distribute the Products and Recommendations within such Member's own organization, whether by manual or electronic means. Each such Member acknowledges and agrees that the foregoing grant is subject to the terms of such Member's membership arrangement with CIS and may, therefore, be modified or terminated by CIS at any time.

Choice of Law; Jurisdiction; Venue
We acknowledge and agree that these Agreed Terms of Use will be governed by and construed in accordance with the laws of the State of Maryland, that any action at law or in equity arising out of or relating to these Agreed Terms of Use shall be filed only in the courts located in the State of Maryland, that we hereby consent and submit to the personal jurisdiction of such courts for the purposes of litigating any such action. If any of these Agreed Terms of Use shall be determined to be unlawful, void, or for any reason unenforceable, then such terms shall be deemed severable and shall not affect the validity and enforceability of any remaining provisions.

Terms of Use Agreement Version 2.1 - 02/20/04

This Page Intentionally Left Blank.

Acknowledgement

We thank the MITRE Corporation for its collaborative effort in the development of this guide. Working closely with our NSA representatives and Center for Internet Security representatives, the MITRE team—Len LaPadula (task leader), Charles Schmidt (technical lead), and Lisa Nordman (task product manager)—generated the security recommendations and rationales in this guide and produced the document.

Trademark Information

This Page Intentionally Left Blank.

Table of Contents

This Page Intentionally Left Blank

List of Figures

List of Tables

This Page Intentionally Left Blank

Introduction

The purpose of this guide is to provide the reader with security configuration guidance for Microsoft's Exchange Server 2003.[1] Furthermore, it is assumed that the underlying operating system is Microsoft's Windows Server 2003. The recommendations contained herein have been tested on a Windows Server 2003-based platform. Although most of the recommendations will apply even if Exchange is loaded over a different Windows OS, no statements regarding security or operability can be made for other platform configurations.

WARNING: This guide does not provide security guidance for the underlying operating system and other functionality (for example, Active Directory, Domain Controllers, etc), except those functions which are specifically necessary for secure implementation and operation of Exchange.

Target Audience

This document is intended for system administrators, but should be read by anyone involved with or interested in installing and/or configuring Exchange. We assume that the reader is a knowledgeable system administrator. In the context of this document, a knowledgeable system administrator is defined as someone who can create and manage accounts and groups, understands how operating systems perform access control, understands how to set account policies and user rights, is familiar with how to set up auditing and read audit logs, and can configure other similar system-related functionality.

Additionally, it is assumed that the reader is a competent Exchange administrator. Consequently, no tutorial-type information is provided regarding Exchange or electronic messaging in general. Many documents and books exist which provide this information, including Microsoft's web presence at http://www.microsoft.com. That site leads to an extensive array of Exchange-related material.

Target Environment

It is critical for the reader to realize that this document is intended to provide security guidance solely for Exchange implementation within a "Specialized-Security - Limited Functionality" environment. This environment is defined as one in which extremely strict security must be implemented due to high mission-criticality or the presence of a significant threat environment. It is expected that few organizations will want to implement *all* recommendations contained in this document, which is why the security level is referred to as "Specialized Security." Additionally, if all recommendations are implemented, the Exchange server will provide only "limited functionality;" a number of necessary messaging functions might not work. It is anticipated that high security environments which these recommendations target will find it necessary to ignore certain recommendations and instead use default settings.

A suggestion to ignore recommended high-security settings may seem counterintuitive to those with a desire for high security, but nearly all readers of this document will find it

[1] Within this document, the term "Exchange" refers to Microsoft's Exchange Server 2003, unless specifically noted otherwise.

necessary to do so in certain situations. For example, an Exchange server is more susceptible to attack if it is running as an NNTP server. Consequently, we recommend that NNTP be disabled. However, if an enterprise's requirements mandate that Exchange be implemented as an NNTP server within a high-security environment, then the recommendation to disable NNTP will need to be disregarded. As a second example, it is imperative to retain logged data for forensic and other purposes. However, an administrator logging large quantities of data may find that disk space quickly becomes full. If sufficient backup space is not available or operational restrictions preclude regular and frequent backup and archival activities, the enterprise might determine that certain high-security logging recommendations are untenable and will need to resort to a lesser security posture.

It is highly recommended that the reader consult related security documents before implementing recommendations contained in this document. For example, *Exchange Server 2003 and Exchange Server Front-End and Back-End Topology* provides security and configuration advice for those implementing Exchange in a "front-end" and "back-end" architecture. When implementing Exchange on the Windows Server 2003 operating system, *Windows Server 2003 Security Guide* should be consulted. If Windows Server 2000 is used, implement security guidance from *Windows Server 2000 Operating System Level 2 Benchmark Consensus Baseline Security Settings*. Information regarding security for Microsoft's Internet Information Services (IIS) can be found at: http://www.microsoft.com/security/guidance/prodtech/IIS.mspx. Finally, for environments with less stringent security requirements such as "legacy" and "enterprise" [2] environments, the reader is referred to *Exchange Server 2003 Security Hardening Guide*.

Getting the Most from this Guide

The following list contains suggestions to facilitate successful use of this guide:

WARNING: This list does not address site-specific issues and every setting or suggestion in this guide should be tested on a non-operational network.

- ❑ Read the guide in its entirety. Omitting or deleting steps can potentially lead to an unstable system and/or network that will require reconfiguration and reinstallation of software.
- ❑ Perform pre-configuration recommendations.
- ❑ Perform a complete backup of your system before implementing any of the recommendations in this guide.
- ❑ Ensure that the latest bug fixes have been installed.
- ❑ Use security settings that are appropriate for your environment.

[2] A "legacy" environment is defined herein as one in which an administrator is migrating in an operational environment from an older Exchange version such as Exchange 5.5 or Exchange Server 2000. An "enterprise" environment is defined herein as one in which Exchange is being implemented within a large organizational unit. Typically, this environment requires stringent security measures, but the large variety of business processes within a big organization require that many Exchange functions be relaxed from their highest possible settings in order to accommodate the large set of messaging functions required.

About the Compliance Checking Tool

The Center for Internet Security offers a compliance checking tool which serves as a companion to this document. Due to technical considerations, not every recommendation in this document is scored by the tool. The table in Chapter 2 identifies which recommendations are checked by the tool, as opposed to those that will require manual inspection. Whether or not the administrator uses the tool to inspect compliance, it will be necessary to manually implement the recommendations. The tool conducts only compliance checking; it does not modify the platform. NOTE: The CIS tool is not currently available. Please check the CIS website at http://www.cisecurity.org for updated information on the tool availability.

About this Guide

This document consists of the following chapters and appendixes:

Chapter 1 – General Exchange Security Guidance

This chapter contains general security guidance for installing and running Exchange. For example, securing the underlying operating system, applying all applicable patches, and ensuring specific services are enabled.

Chapter 2 – Recommendation Security Settings for Exchange Controls

Chapter 2 provides security recommendations for specific controls and describes how to manually implement the recommendations.

Chapter 3 – Rationales for Exchange Control Security Settings

Chapter 3 supplies rationales for the security recommendations made in Chapter 2. Rationales provide a brief description of the control's functionality, recommended value, and a synopsis of the risks if the recommended value is not implemented.

Appendix A – Bibliography

This Page Intentionally Left Blank.

Chapter

1

General Exchange Security Guidance

This chapter contains general security guidance for Exchange Server 2003. The following recommendations are provided to facilitate a more secure platform.

- ❑ Review all recommendations to ensure they comply with local policy.
- ❑ Do not install Exchange Server 2003 on a domain controller.
- ❑ Load the operating system and secure it before loading Exchange onto the platform. It is important to realize that the system cannot be considered to be secured until the operating system has first been secured. If the operating system is not secured, Exchange functionality *might* be secure but the platform as a whole will be vulnerable.
- ❑ Ensure the following services have been started before attempting to install Exchange:
  - ❑ NNTP
  - ❑ HTTP
  - ❑ SMTP
  - ❑ World Wide Web
  - ❑ .NET Framework
- ❑ Ensure that all relevant operating system security patches have been applied.
- ❑ Ensure that all relevant Exchange security patches have been applied.
- ❑ Exchange Administrator should require a User's network/domain username to be different than their email alias. The possible threat in not following this recommendation: once a malicious user has access to your email address, they now have a valid network/domain username to conduct malicious activity.
- ❑ Subscribe to the Windows Security mailing list at: http://www.microsoft.com/technet/security/bulletin/notify.mspx
- ❑ Visit and implement appropriate recommendations at the vendor's web site: http://www.microsoft.com/exchange/techinfo/security/default.asp
- ❑ Because the recommendations in this document yield a Specialized Security - Limited Functionality posture, it is essential that the administrator:
  - ❑ Carefully review the recommendations to ensure that they are not disabling functionality that is required by their enterprise.
  - ❑ Test the settings in a non-production environment before deployment.
- ❑ The recommended settings only *increase* security. It is essential to continually monitor the latest in best security practices.

Native Mode vs. Mixed Mode

Exchange Server 2003 can operate in two modes: native mode and mixed mode. Exchange 2003 servers running in native mode can rename and consolidate administrative groups, define routing groups and administrative groups, move mailboxes between servers in different administrative groups, create an administrative group that spans multiple routing groups, and use query-based distribution groups. Mixed mode environments do not provide that functionality. Native mode does not allow Exchange 2003 to interoperate with Exchange 5.5 systems.

Mixed mode was designed for Exchange 2003 to interoperate with Exchange 5.5 systems, and is the default mode. If your environment contains Exchange 5.5 systems, mixed mode should be used. To switch from mixed mode to native mode, all of the Exchange servers in your organization must be running Exchange 2003 or Exchange 2000. Once the Exchange servers have been updated, the switch to Native mode can take place. Once the switch occurs, the change cannot be reversed, and the organization is no longer able to interoperate with Exchange 5.5 systems. For more information about how to switch to native mode, see the *Exchange Server 2003 Deployment Guide* (http://go.microsoft.com/fwlink/?linkid=14576).

Front-End/Back-End Server Configuration

Exchange 2003 supports front-end and back-end server architectures. The front-end server is a server running Exchange 2003 that does not host mailboxes or public folder stores. The front-end server forwards client requests to the back-end server for processing. Mailboxes and/or public folder stores are hosted on the back-end server(s) that are running Exchange 2003. The advantages of using front-end/back-end server topology are as follows:

- A single namespace for each user.
- Offload processing.
- Strengthened security: Front-end server can be positioned as the single point of access and located after a firewall(s).
- Scalability: As an enterprise expands, adding additional front-end and back-end servers can assist in the expansion as well as address load balancing issues.

Figure 1 depicts the recommended Exchange front-end and back-end server architecture. It should be noted that while we recommend placing a firewall in front of the front-end server, we were unable to configure this in our small lab environment. In a real world environment, it is critical to place a firewall in front of each front-end machine. This adds an additional layer of protection between the internet and the front-end server.

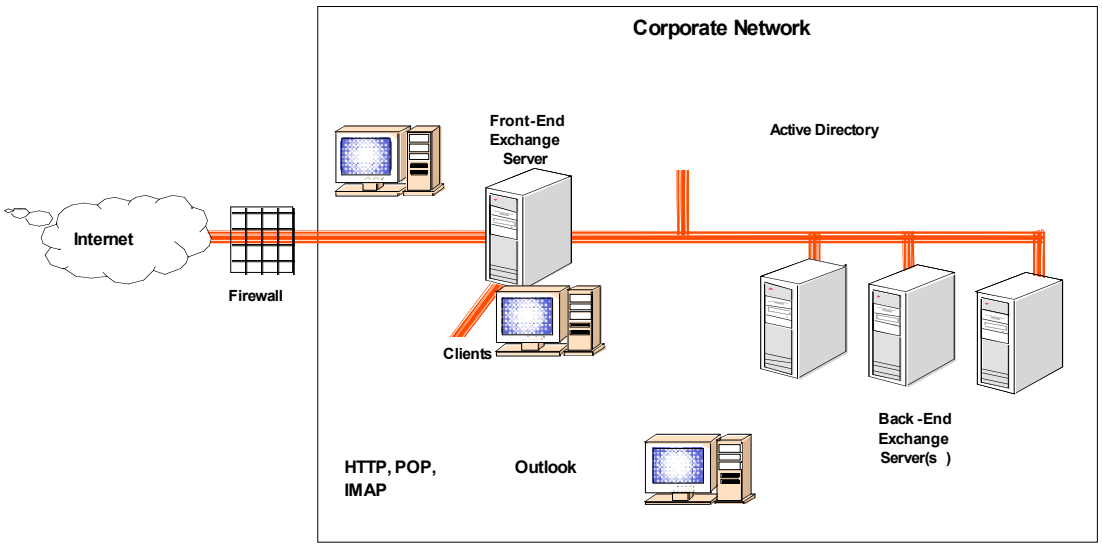

Figure 1. Recommended Exchange Front-end/Back-end Server Architecture

For additional information on how to set up Exchange front-end and back-end architecture, refer to Microsoft's Exchange Server 2003 and Exchange 2000 Server Front-End and Back-End Topology (http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/febetop.mspx).

Roles

Exchange 2003 Servers can be configured to function with a specific role. That is, as an HTTP Server, IMAP4 Server, POP3 Server, NNTP Server or SMTP Server. This paper provides security recommendations based on each of these roles. Table 1 identifies services on the server that must be enabled and started for a role to function. For example, if your server was configured as a POP3 server, both Microsoft Exchange POP3 service and the IIS Manager Service should be started. The SMTP service must be running on every Exchange Server 2003, because the entire Exchange Server 2003 messaging system depends on it. Without the SMTP service, Exchange will not function.

Table 1. Roles and Services That Must Be Enabled and Started

	Required Services	Comments
HTTP Role	World Wide Web (WWW) publishing service	Server used for HTTP
	HTTP SSL	Starts when required by WWW
	IIS Admin Service	Required if running WWW, SMTP, POP3, IMAP4, or NNTP services
POP3 Role	Microsoft Exchange POP3	Server used for POP3
	IIS Admin Service	Required if running WWW, SMTP, POP3, IMAP4, or NNTP services
IMAP4 Role	Microsoft Exchange IMAP4	Server used for IMAP4
	IIS Admin Service	Required if running WWW, SMTP, POP3, IMAP4, or NNTP services
SMTP Role	Simple Mail Transport Protocol (SMTP)	Server used for SMTP and required service for Exchange to function
	IIS Admin Service	Required if running WWW, SMTP, POP3, IMAP4, or NNTP services
	Microsoft Exchange Information Store	Used by virus scanners, SMTP
	Microsoft Exchange System Attendant	Required for Exchange maintenance
	Microsoft Exchange Management	Required for message tracking functionality
	Microsoft Exchange MTA Stacks	Used for error handling of some messages
NNTP Role	Network News Transport Protocol	Server used for NNTP
	IIS Admin Service	Required if running WWW, SMTP, POP3, IMAP4, or NNTP services

Logs

There are a number of logging capabilities in Exchange 2003. For example, diagnostic logging, message tracking, activity logging, and audit logging. The volume of data these logs can generate in a day can be enormous and cause Exchange to fail. Exchange will dismount its stores if it detects that it has run out of disk space, resulting in a complete loss of Exchange services. To avoid potential problems with logging data, it is required that all log files be written to separate partitions from the partitions used by the Exchange Stores.

Exchange Features

Outlook Mobile Access (OMA): OMA is a text interface for accessing Exchange server based email, calendars, contacts, and tasks. It was designed to be accessed with PDAs and/or telephones with small displays and web based capabilities. By default OMA is disabled. We recommend keeping the default at "disabled" unless there is specific enterprise policy/need to enable OMA.

ActiveSync: ActiveSync is a service provided by Exchange Server 2003 that allows users to synchronize e-mail, calendaring, and contact information between the Exchange server and Windows supporting mobile devices, such as PDAs. When Exchange is installed, ActiveSync is enabled by default. Disable ActiveSync, unless enterprise policy states it is needed. Disabling ActiveSync restricts connectivity methods to the server, helps to reduce attack vectors as well as reduce the ways in which viruses may be introduced to your network, and provides (admittedly limited) control over how easily messages and other posts are allowed to leave your exchange network.

Outlook Web Access (OWA): OWA is a remote email client and a basic Exchange Server 2003 messaging service. OWA allows users to access mailboxes and folders using an Internet browser that is compliant with HTML 3.2 and JavaScript standards. The HTML standard provides drag and drop functionality, expandable folder hierarchies, HTML composition, toolbar tips and Kerberos authentication. The newest OWA client does not support offline mode, Outlook journaling, or copying or moving items between folders. OWA is installed automatically with Exchange Server 2003, and cannot be removed. Disable OWA, unless enterprise policy states it is needed. Disabling OWA restricts connectivity methods to the server, helps reduce attack vectors, and reduce the ways in which viruses maybe introduced to the network.

Public Folders: Exchange Server 2003 supports public folders. The purpose of public folders is to provide common access to messages and files. Files can be dragged and dropped into public folders for instant sharing. In addition, sorting rules can be applied to public folders to ensure items are arranged by name or date. Each server has one default public folder store (named the **Public Folder Store**) that supports the **Public Folder** tree. There can be up to five stores per storage group, and any number of the five stores may be public folders. Lastly, users can be granted access or denied access to specific public folders.

Public Folder Tree: Public Folder Trees are a group of public folders in a hierarchical structure. One tree can have multiple public folder stores. Each Exchange organization has one default public tree. Additional public folder trees can be created that users can access using Outlook Web Access.

Mailbox Stores: Mailbox Stores is a database for storing mailboxes in Exchange 2003. Mailbox stores hold data that is private to an individual and contain mailbox folders generated each time a new mailbox is created for an individual. Stores are located in storage groups. There can be up to five stores per storage group, and any number of the five stores may be mailbox stores. Each mailbox store has a set of transaction logs associated with it. These transaction logs provide detailed information of messages received and sent from a store in a storage group.

Important Security Points

- Do not install Exchange Server 2003 on a domain controller.
- Ensure that all relevant operating system security patches have been applied.
- Ensure that all relevant Exchange security patches have been applied.
- Exchange Administrator should require a User's network/domain username to be different than their email alias. The possible threat in not following this recommendation: once a malicious user has access to your email address, they now have a valid network/domain username to conduct malicious activity.
- Because the recommendations in this document yield a Specialized Security - Limited Functionality posture, it is essential that the administrator:
  - Carefully review the recommendations to ensure that they are not disabling functionality that is required by their enterprise.
- Strengthen security by using Front-end/Back-end server configuration.
- Configure log files so that they are written to partitions that are separate from the partitions used by the Exchange Stores.

Chapter

2

Recommended Security Settings for Exchange Controls

This chapter supplies control-specific Exchange guidance. (In this document, the term "control" refers to a specific parameter or setting through which the administrator configures the platform through the user-interface.)

Table 2 identifies the specific security recommendations. The information contained in the control table is the primary means by which the administrator can secure the Exchange platform. Several types of information are provided in the control table. The fundamental security-related information provided includes:

- The control name
- The recommended security setting for the control
- How the administrator can navigate through the user-interface to access that control

The following is an explanation of the column and row headings in the table.

Column Headings:

- Reference ID/Control Title—Number and Title associated with specific control for referencing purposes.
- Roles: HTTP Server—Recommended values for Exchange Server platforms which have been implemented as HTTP servers.
- Roles: POP3 Server—Recommended values for Exchange Server platforms which have been implemented as POP3 servers.
- Roles: IMAP4 Server—Recommended values for Exchange Server platforms which have been implemented as IMAP4 servers.
- Roles: NNTP Server—Recommended values for Exchange Server platforms which have been implemented as NNTP servers.
- Roles: SMTP Server—Recommended values for Exchange Server platforms which have been implemented as SMTP servers.
- S/R/Manual—Identifies whether a control is included in the compliance-checking tool assessment-of-compliance ("scorable"), or whether the control value is only reported by the tool and not included in the assessment ("reportable"). Lastly, items that are marked ("manual"), means the tool can neither score nor report on the item.
- GUI Location and Control Name in Exchange—Navigation path through the Exchange GUI to locate the control setting and the name of the control as it appears in the MS Exchange Graphical User Interface (GUI). The GUI path assumes that the display administrative groups and display routing groups have been checked.

- Description—A short description of the capability provided by the control.
- Default—The controls initial value upon system initialization.

Row Headings:

- **Global Settings**: Security controls directly related to global settings.
- **HTTP**: Security controls directly related to the HTTP protocol.
- **IMAP**: Security controls directly related to the IMAP4 protocol.
- **Mailbox Stores**: Security controls directly related to mailbox stores.
- **NNTP**: Security controls directly related to network news transport protocol.
- **POP3**: Security controls directly related to the POP3 protocol.
- **Public Folder Stores**: Security controls directly related to public folder stores.
- **Public Folders**: Security controls directly related to public folders.
- **Routing Groups**: Security controls directly related to routing groups.
- **Servers**: Security controls directly related to servers.
- **SMTP**: Security controls directly related to the SMTP protocol.
- **User Level Controls**: Security controls directly related to user level controls.

Row Sub-headings:

- **Admin Post Requests**: Items related to deleting or posting specific articles, and creating or deleting newsgroups.
- **Authentication**: Items related to authentication options.
- **Automated Messages**: Items related to automatic response, forwarding, and delivery messages.
- **Backup/Restore**: Items related to backup and restores. For example, last time of full/incremental backup, do not delete items until fully backed up, and maintenance intervals.
- **Connections**: Items related to connection limits and timeouts, SSL connection options, port and relay.
- **Deletion Settings**: Items related to deletion settings, i.e., keep deleted items for x-number of days and age limits on public folders controls.
- **Delivery Restrictions**: Items related to filtering, recipient limits, message size, user and list restrictions.

- **Directory Access**: Items related to accessing directories, i.e., scripts, executables, and access controls.
- **Display**: Items which display policies.
- **Enabled Services**: Items related to enabling services.
- **Logging**: Items related to enabling logging, audit tracking, and archiving controls.
- **Mounting Stores**: Items related to mounting or not-mounting specific stores at startup.
- **Permissions**: Items related to group and/or username permissions, permissions of administration of public folders, permissions of exchange recipients and roles.
- **Storage Limits**: Items related to storage limits.
- **Visibility**: Items that control whether or not a public store was visible.

Sometimes there is more than one way through the user-interface to implement a setting. It is highly recommended to configure the Exchange GUI to display both *administrative groups* and *routing groups*. This configuration directly corresponds to the configuration used by this document when it describes object container hierarchies for various controls. Failure to follow this recommendation will not make your Exchange server any less secure, but it may make this document slightly more difficult to use. In most cases it is irrelevant as to which path or functionality is used. However, in cases where the method will make a difference in the final security posture, these are identified. For example, settings made from the IIS interface are usually overwritten by settings made through the Exchange System Manager (ESM) interface. Because of this, the rationales for these controls direct the administrator to configure the security setting using ESM.

Table 2. Security Recommendations

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
1. Global Settings									
1.1 Automated Messages									
1.1.1 Allow out-of-office Responses	Clear all checkboxes	Clear all checkboxes	Clear all checkboxes	Clear all checkboxes	Clear all checkboxes	S	Exchange System Manager → Global settings → Internet Message Formats → [specific format] → Properties → Advanced Tab → (Multiple Items)	Automatic message creation controls.	Allow "Non-delivery reports" is selected.
1.2 Delivery Restrictions									
1.2.1 Global Accept and Deny List Configuration	NR	NR	NR	NR	NR	R	Exchange System Manager → Global Settings → Message Delivery → Properties → Connection Filtering Tab → Global Accept and Deny List Configuration (Multiple Items)	Global Accept and Deny List Configuration.	None
1.2.2 Specify Block List Service Provider	Enter the display name, DNS suffix of the provider, and a return message to return.	Enter the display name, DNS suffix of the provider, and a return message to return.	Enter the display name, DNS suffix of the provider, and a return message to return.	Enter the display name, DNS suffix of the provider, and a return message to return.	Enter the display name, DNS suffix of the provider, and a return message to return.	S	Exchange System Manager → Global Settings → Message Delivery → Properties → Connection Filtering Tab → Block List Service Configuration → Add Button	Specify a block list/blacklist provider to use for spam filtering.	None
1.2.3 Block List Exceptions	NR	NR	NR	NR	NR	R	Exchange System Manager → Global Settings → Message Delivery → Properties → Connection Filtering Tab → Block List Service Configuration → Exception Button	Add SMTP addresses that should not be blocked despite being in blacklist.	None
1.2.4 Sending/Receiving Message Size	<=30 MB	<=30 MB	<=30 MB	<=30 MB	<=30 MB	S	Exchange System Manager → Global Settings → Message Delivery → Properties → Defaults Tab → Sending message size and Receiving message size	Sending|receiving message size.	Both limits set at 10MB.
1.2.5 Recipient Limits	<=5,000	<=5,000	<=5,000	<=5,000	<=5,000	S	Exchange System Manager → Global Settings → Message Delivery → Properties → Defaults Tab → Recipient limits	Recipient limits (max number of recipients for a single message from this server).	Limit set to 5000.
1.2.6 Filter Recipients who are not in Directory	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Global Settings → Message Delivery → Properties → Recipient Filtering Tab → Filter recipients who are not in the Directory	Filter recipients who are not in the Directory (block any recipient not listed in domain's AD).	Cleared
1.2.7 Recipients	NR	NR	NR	NR	NR	R	Exchange System Manager → Global Settings → Message Delivery → Properties → Recipient Filtering	Block messages that are sent to the following recipients.	None

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
							Tab → Recipients		
1.2.8 Senders	NR	NR	NR	NR	NR	R	Exchange System Manager → Global Settings → Message Delivery → Properties → Sender Filtering Tab → Senders	Block messages that claim to be from the following senders.	None
1.2.9 Archive Filtered Messages	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Global Settings → Message Delivery → Properties → Sender Filtering Tab → Archive filtered messages	Archive filtered messages.	Cleared
1.2.10 Filter Messages with Blank Sender	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Global Settings → Message Delivery → Properties → Sender Filtering Tab → Filter messages with blank sender	Filter messages with blank sender.	Cleared
1.2.11 Drop Connection if Address Matches Filter	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Global Settings → Message Delivery → Properties → Sender Filtering Tab → Drop connection if address matches filter	Drop connection if address matches filter (immediately drop the connection if message would be filtered).	Selected
1.2.12 Accept Messages Without Notifying Sender of Filtering	Select the checkbox (cannot be set if "Drop connection if address matches filter" [item 1.2.11] is selected)	Select the checkbox (cannot be set if "Drop connection if address matches filter" [item 1.2.11] is selected)	Select the checkbox (cannot be set if "Drop connection if address matches filter" [item 1.2.11] is selected)	Select the checkbox (cannot be set if "Drop connection if address matches filter" [item 1.2.11] is selected)	Select the checkbox (cannot be set if "Drop connection if address matches filter" [item 1.2.11] is selected)	S	Exchange System Manager → Global Settings → Message Delivery → Properties → Sender Filtering Tab → Accept messages without notifying sender of filtering	Accept messages without notifying sender of filtering. (Filter silently to avoid giving indications that messages were filtered to the sender.)	Cleared
1.3 Enabled Services									
1.3.1 Enable Outlook Mobile Access	Clear the checkbox and delete the OMA Virtual Directory from IIS	Clear the checkbox and delete the OMA Virtual Directory from IIS	Clear the checkbox and delete the OMA Virtual Directory from IIS	Clear the checkbox and delete the OMA Virtual Directory from IIS	Clear the checkbox and delete the OMA Virtual Directory from IIS	S	Exchange System Manager → Global Settings → Mobile Services → Properties → General Tab → Outlook Mobile Access → Enable Outlook Mobile Access and Enable unsupported devices	Enable Outlook Mobile access and Enable unsupported devices (do not require recipient device to be recognized).	Both cleared.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
1.3.2 Enable User Initiated Synchronization	Clear the checkbox and delete the Microsoft-Server-ActiveSync Virtual Directory	Clear the checkbox and delete the Microsoft-Server-ActiveSync Virtual Directory	Clear the checkbox and delete the Microsoft-Server-ActiveSync Virtual Directory	Clear the checkbox and delete the Microsoft-Server-ActiveSync Virtual Directory	Clear the checkbox and delete the Microsoft-Server-ActiveSync Virtual Directory	S	Exchange System Manager → Global Settings → Mobile Services → Properties → General Tab → Exchange ActiveSync → Enable user initiated synchronization/Enable up-to-date notifications/Enable notifications to user specified SMTP addresses	Enable user initiated synchronization (if disabled, turns off ActiveSync), Enable up-to-date notifications (alerts from the server that handheld should resync) and Enable notifications to user specified SMTP addresses.	All three cleared.
1.4 Logging									
1.4.1 Archive Filtered Messages							See item 1.2.9	See item 1.2.9	
2. HTTP									
2.1 Authentication									
2.1.1 Certificate Wizard	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	S	IIS Manager → [server] → Web Sites → Default Web Site → Properties → Directory Security Tab → Server Certificate Button → Wizard Button	Install or modify certificates.	When clicked, the web server certificate wizard runs.
2.1.2 Enable Anonymous Access	NR	NR	NR	NR	NR	R	IIS Manager → [server] → Web Sites → Default Web Site → Properties → Directory Security Tab → Authentication and access control → Edit Button → Enable anonymous access	Whether access to the HTTP service should be anonymous.	Cleared
2.1.3 Username/Password	Use the IUSR_<computer-name> account	Use the IUSR_<computer-name> account	Use the IUSR_<computer-name> account	Use the IUSR_<computer-name> account	Use the IUSR_<computer-name> account	R	IIS Manager → [server] → Web Sites → Default Web Site → Properties → Directory Security Tab → Authentication and access control → Edit Button → Enable anonymous access → User name and Password	When users connect to the HTTP server anonymously, what user name they should they use.	IUSR_<computer *name*> & associated password.
2.1.4 Authenticated Access	Select integrated windows authentication checkbox	Select integrated windows authentication checkbox	Select integrated windows authentication checkbox	Select integrated windows authentication checkbox	Select integrated windows authentication checkbox	S	IIS Manager → [server] → Web Sites → Default Web Site → Properties → Directory Security Tab → Authentication and access control → Edit Button → Authenticated access (Multiple Items)	Which authentication methods should be used to connect to the HTTP server.	Integrated Windows authentication is selected.
2.1.5 Authentication Method to Access Exchange Virtual Directory	Delete the Exchange Virtual Directory or (if Outlook Web Access must be used), enable only	Delete the Exchange Virtual Directory or (if Outlook Web Access must be used), enable only	Delete the Exchange Virtual Directory or (if Outlook Web Access must be used), enable only	Delete the Exchange Virtual Directory or (if Outlook Web Access must be used), enable only	Delete the Exchange Virtual Directory or (if Outlook Web Access must be used), enable only	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → HTTP → Exchange Virtual Server → Exchange → Properties → Access Tab → Authentication Settings → Authentication Button	One or more methods of authentication must be selected.	Integrated Windows Authentication and Basic authentication selected.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
	Integrated Windows Authentication and Basic authentication.	Integrated Windows Authentication and Basic authentication.	Integrated Windows Authentication and Basic authentication.	Integrated Windows Authentication and Basic authentication.	Integrated Windows Authentication and Basic authentication.				
2.1.6 Authentication Method to Access Exadmin Virtual Directory	Integrated Windows authentication selected	Integrated Windows authentication selected	Integrated Windows authentication selected	Integrated Windows authentication selected	Integrated Windows authentication selected	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → HTTP → Exchange Virtual Server → Exadmin → Properties → Access Tab → Authentication Settings → Authentication Button	One or more methods of authentication must be selected.	Integrated Windows Authentication selected.
2.1.7 Authentication Method to Access ActiveSync Virtual Directory	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) select only basic authentication	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) select only basic authentication	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) select only basic authentication	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) select only basic authentication	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) select only basic authentication	S	IIS Manager → [server] → Web Sites → Default Web Site → Microsoft Server-ActiveSync → Properties → Directory Security Tab → Authentication and access control → Edit Button → Authenticated access (Multiple Items)	One or more methods of authentication must be selected.	Basic authentication
2.1.8 Authentication Method to Access Outlook Mobile Access Virtual Directory	Delete the OMA Virtual Directory or (if OMA is used) select only basic authentication	Delete the OMA Virtual Directory or (if OMA is used) select only basic authentication	Delete the OMA Virtual Directory or (if OMA is used) select only basic authentication	Delete the OMA Virtual Directory or (if OMA is used) select only basic authentication	Delete the OMA Virtual Directory or (if OMA is used) select only basic authentication	S	IIS Manager → [server] → Web Sites → Default Web Site → OMA → Properties → Directory Security Tab → Authentication and access control → Edit Button → Authenticated access (Multiple Items)	One or more methods of authentication must be selected.	Basic authentication
2.1.9 Enable Forms Based Authentication	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → HTTP → Exchange Virtual Server → Properties → Settings Tab → Outlook Web Access → Enable Forms Based Authentication	Determine whether web forms can provide authentication information.	Enable Forms Based Authentication checkbox is cleared.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
2.1.10 Authentication Method to Access Public Virtual Directory	Delete the Public Virtual Directory or (if public folders are used, select only Integrated Windows Authentication and Basic authentication.	Delete the Public Virtual Directory or (if public folders are used, select only Integrated Windows Authentication and Basic authentication.	Delete the Public Virtual Directory or (if public folders are used, select only Integrated Windows Authentication and Basic authentication.	Delete the Public Virtual Directory or (if public folders are used, select only Integrated Windows Authentication and Basic authentication.	Delete the Public Virtual Directory or (if public folders are used, select only Integrated Windows Authentication and Basic authentication.	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → HTTP → Exchange Virtual Server → Public → Properties → Access Tab → Authentication Settings → Authentication Button	One or more methods of authentication must be selected.	Integrated Windows authentication and Basic authentication are selected.
2.1.11 Require Secure Channel/Require 128 bit encryption to Microsoft-Server-ActiveSync Virtual Directory	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) require both a secure channel and 128 bit encryption.	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) require both a secure channel and 128 bit encryption.	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) require both a secure channel and 128 bit encryption.	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) require both a secure channel and 128 bit encryption.	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) require both a secure channel and 128 bit encryption.	S	IIS Manager → [server] → Web Sites → Default Web Site → Microsoft Server ActiveSync → Properties → Directory Security Tab → Secure communications → Edit Button → Require secure channel (SSL) (Multiple Items)	Server Communications Options.	Cleared
2.1.12 Require Secure Channel (SSL) and Require 128 bit encryption to OMA Virtual Directory	Delete the OMA Virtual Directory or (if OMA is used) require both a secure channel and 128 bit encryption.	Delete the OMA Virtual Directory or (if OMA is used) require both a secure channel and 128 bit encryption.	Delete the OMA Virtual Directory or (if OMA is used) require both a secure channel and 128 bit encryption.	Delete the OMA Virtual Directory or (if OMA is used) require both a secure channel and 128 bit encryption.	Delete the OMA Virtual Directory or (if OMA is used) require both a secure channel and 128 bit encryption.	S	IIS Manager → [server] → Web Sites → Default Web Site → OMA → Properties → Directory Security Tab → Secure communications → Edit Button → Require secure channel (SSL) (Multiple Items)	Server Communications Options.	Cleared

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
2.1.13 Require Secure Channel (SSL) and Require 128 bit encryption to Public Virtual Directory	Delete the Public Virtual Directory or (if public folders are used) require both a secure channel and 128 bit encryption.	Delete the Public Virtual Directory or (if public folders are used) require both a secure channel and 128 bit encryption.	Delete the Public Virtual Directory or (if public folders are used) require both a secure channel and 128 bit encryption.	Delete the Public Virtual Directory or (if public folders are used) require both a secure channel and 128 bit encryption.	Delete the Public Virtual Directory or (if public folders are used) require both a secure channel and 128 bit encryption.	S	IIS Manager → [server] → Web Sites → Default Web Site → Public → Properties → Directory Security Tab → Secure communications → Edit Button → Require secure channel (SSL) (Multiple Items)	Server Communications Options.	Cleared
2.1.14 Require Secure Channel (SSL) and Require 128 bit encryption to Exchange Virtual Directory	Delete the Exchange Virtual Directory or (if OWA is used) require both a secure channel and 128 bit encryption.	Delete the Exchange Virtual Directory or (if OWA is used) require both a secure channel and 128 bit encryption.	Delete the Exchange Virtual Directory or (if OWA is used) require both a secure channel and 128 bit encryption.	Delete the Exchange Virtual Directory or (if OWA is used) require both a secure channel and 128 bit encryption.	Delete the Exchange Virtual Directory or (if OWA is used) require both a secure channel and 128 bit encryption.	S	IIS Manager → [server] → Web Sites → Default Web Site → Exchange → Properties → Directory Security Tab → Secure communications → Edit Button → Require secure channel (SSL) (Multiple Items)	Server Communications Options.	Cleared
2.1.15 Require Secure Channel (SSL) and Require 128 bit encryption to Exadmin Virtual Directory	Clear the checkboxes	Clear the checkboxes	Clear the checkboxes	Clear the checkboxes	Clear the checkboxes	S	IIS Manager → [server] → Web Sites → Default Web Site → Exadmin → Properties → Directory Security Tab → Secure communications → Edit Button → Require secure channel (SSL) (Multiple Items)	Server Communications Options.	Cleared
2.2 Connections									
2.2.1 TCP Port/SSL Port	80/443	80/443	80/443	80/443	80/443	S	IIS Manager → [server] → Web Sites → Default Web Site → Properties → Web Site Tab → Web site identification → TCP port and SSL port	TCP port used by HTTP and HTTPS daemon.	80/443
2.3 Delivery Restrictions									
2.3.1 Address/Domain Name Restrictions	NR	NR	NR	NR	NR	R	IIS Manager → [server] → Web Sites → Default Web Site → Properties → Directory Security Tab → IP Address and domain name restrictions → Edit Button	Whether access is granted or denied by default and any IP addresses (or subnets) that should be exceptions to the general rule.	Default granted to all with no exceptions.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
2.4 Directory Access									
2.4.1 Execute Permissions for OWA Virtual Directory	Delete the Exchange Virtual Directory or (if OWA is used) select none	Delete the Exchange Virtual Directory or (if OWA is used) select none	Delete the Exchange Virtual Directory or (if OWA is used) select none	Delete the Exchange Virtual Directory or (if OWA is used) select none	Delete the Exchange Virtual Directory or (if OWA is used) select none	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → HTTP → Exchange Virtual Server → Exchange → Properties → Access Tab → Execute permissions	Whether scripts, scripts and executables, or neither can be executed.	"None" is selected [other options are scripts, scripts and executables.
2.4.2 Access Control to OWA Virtual Directory	Delete the Exchange Virtual Directory or (if OWA is used) select read, write, script source access, and directory browsing	Delete the Exchange Virtual Directory or (if OWA is used) select read, write, script source access, and directory browsing	Delete the Exchange Virtual Directory or (if OWA is used) select read, write, script source access, and directory browsing	Delete the Exchange Virtual Directory or (if OWA is used) select read, write, script source access, and directory browsing	Delete the Exchange Virtual Directory or (if OWA is used) select read, write, script source access, and directory browsing	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → HTTP → Exchange Virtual Server → Exchange → Properties → Access Tab → Access Control	Read, write, script source access, or directory browsing.	Read, write, script source access, directory browsing are selected.
2.4.3 Access Controls to Exadmin Virtual Directory	Select read, write, script source access, and directory browsing	Select read, write, script source access, and directory browsing	Select read, write, script source access, and directory browsing	Select read, write, script source access, and directory browsing	Select read, write, script source access, and directory browsing	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → HTTP → Exchange Virtual Server → Exadmin → Properties → Access Tab → Access Control (Multiple Items)	Read, write, script source access, or directory browsing.	Read, write, script source access, directory browsing are selected.
2.4.4 Execute Permissions to access Exadmin Virtual Directory	Select "None"	Select "None"	Select "None"	Select "None"	Select "None"	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → HTTP → Exchange Virtual Server → Exadmin → Properties → Access Tab → Execute Permissions (Multiple Items)	Whether scripts, scripts and executables, or neither can be executed.	None
2.4.5 Access Control to ActiveSync Virtual Directory	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) grant read access only	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) grant read access only	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) grant read access only	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) grant read access only	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) grant read access only	S	IIS Manager → [server] → Web Sites → Default Web Site → Microsoft-Server-ActiveSync → Properties → Virtual Directory Tab → Access Control (Multiple Items)	Read, write, script source access, or directory browsing.	Read is enabled. (This cannot be changed in this GUI.)

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
2.4.6 Execute Permissions to access ActiveSync Virtual Directory	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) grant scripts and executables the ability to run	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) grant scripts and executables the ability to run	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) grant scripts and executables the ability to run	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) grant scripts and executables the ability to run	Delete the ActiveSync Virtual Directory or (if ActiveSync is used) grant scripts and executables the ability to run	S	IIS Manager → [server] → Web Sites → Default Web Site → Microsoft-Server-ActiveSync → Properties → Virtual Directory Tab → Execute permissions (Multiple Items)	Whether scripts, scripts and executables, or neither can be executed.	Scripts and Executables are selected. (This cannot be changed in this GUI.)
2.4.7 Access Control to OMA Virtual Directory	Delete OMA Virtual Directory or (if OMA is to be used) grant read access only	Delete OMA Virtual Directory or (if OMA is to be used) grant read access only	Delete OMA Virtual Directory or (if OMA is to be used) grant read access only	Delete OMA Virtual Directory or (if OMA is to be used) grant read access only	Delete OMA Virtual Directory or (if OMA is to be used) grant read access only	S	IIS Manager → [server] → Web Sites → Default Web Site → OMA → Properties → Virtual Directory Tab → Access Control (Multiple Items)	Read, write, script source access, or directory browsing.	Read is enabled. (This cannot be changed in this GUI.)
2.4.8 Execute Permissions to access OMA Virtual Directory	Delete OMA Virtual Directory of (if OMA is to be used) enable scripts only to be run	Delete OMA Virtual Directory of (if OMA is to be used) enable scripts only to be run	Delete OMA Virtual Directory of (if OMA is to be used) enable scripts only to be run	Delete OMA Virtual Directory of (if OMA is to be used) enable scripts only to be run	Delete OMA Virtual Directory of (if OMA is to be used) enable scripts only to be run	S	IIS Manager → [server] → Web Sites → Default Web Site → OMA → Properties → Virtual Directory Tab → Execute permissions (Multiple Items)	Whether scripts, scripts and executables, or neither can be executed.	Scripts are enabled. (This cannot be changed in this GUI.)
2.4.9 Execute Permissions to access Public Virtual Directory	Delete the Public Virtual Directory of (if public folders are to be used) select none.	Delete the Public Virtual Directory of (if public folders are to be used) select none.	Delete the Public Virtual Directory of (if public folders are to be used) select none.	Delete the Public Virtual Directory of (if public folders are to be used) select none.	Delete the Public Virtual Directory of (if public folders are to be used) select none.	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → HTTP → Exchange Virtual Server → Public → Properties → Access Tab → Execute Permissions	Whether scripts, scripts and executables, or neither can be executed.	None is selected [other options are scripts, scripts and executables.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
2.4.10 Access Controls to Public Virtual Directory	Delete the Public Virtual Directory or (if public folders are to be used) select read, write, script source access and directory browsing.	Delete the Public Virtual Directory or (if public folders are to be used) select read, write, script source access and directory browsing.	Delete the Public Virtual Directory or (if public folders are to be used) select read, write, script source access and directory browsing.	Delete the Public Virtual Directory or (if public folders are to be used) select read, write, script source access and directory browsing.	Delete the Public Virtual Directory or (if public folders are to be used) select read, write, script source access and directory browsing.	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → HTTP → Exchange Virtual Server → Public → Properties → Access Tab → Access Control	Read, write, script source access, or directory browsing.	Read, write, script source access, directory browsing are selected.
2.5 Logging									
2.5.1 Enable Logging	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	IIS Manager → [server] → Web Sites → Default Web Site → Properties → Web Site Tab → Enable Logging	Whether user activity should be logged.	Selected
3. IMAP									
3.1 Authentication									
3.1.1 Authentication Method used to Access IMAP Virtual Directory	Select the "Basic authentication" and "SSL/TLS" checkboxes	Select the "Basic authentication" and "SSL/TLS" checkboxes	Select the "Basic authentication" and "SSL/TLS" checkboxes	Select the "Basic authentication" and "SSL/TLS" checkboxes	Select the "Basic authentication" and "SSL/TLS" checkboxes	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → IMAP4 → [Specific IMAP4 Virtual Server] → Properties → Access Tab → Access control → Authentication Button	Authentication (which authentication methods should be enabled).	Basic and simple (NTLM).
3.1.2 Certificate Wizard	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → IMAP4 → [Specific IMAP4 Virtual Server] → Properties → Access Tab → Certificate Button → Wizard Button	Opens a wizard to create and install certificate.	When clicked, the web server certificate wizard runs.
3.2 Connections									
3.2.1 Require Secure Channel (SSL) and Require 128 bit encryption to IMAP Virtual Directory	Select both checkboxes	Select both checkboxes	Select both checkboxes	Select both checkboxes	Select both checkboxes	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → IMAP4 → [Specific IMAP4 Virtual Server] → Properties → Access Tab → Secure communication → Communication Button → Require Secure Channel and Require 128-bit encryption	Require a secure channel (between IMAP servers).	Not available until certificate is added.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
3.2.2 Use SSL Connections	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → IMAP4 → [Specific IMAP4 Virtual Server] → Properties → Calendaring Tab → Use SSL connections	Use of SSL when downloading meeting requests.	Cleared
3.2.3 TCP Port/SSL Port	Enter 143/993, respectively	Enter 143/993, respectively	Enter 143/993, respectively	Enter 143/993, respectively	Enter 143/993, respectively	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → IMAP4 → [Specific IMAP4 Virtual Server] → Properties → General Tab → Advanced Button → Edit Button → TCP port and SSL port	Specify which ports are used for normal and SSL connections.	143/993 respectively.
3.2.4 Limit Number of Connections	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → IMAP4 → [Specific IMAP4 Virtual Server] → Properties → General Tab → Limit number of connections to	Limit number of connections to: (maximum number of simultaneous connections to the virtual server).	Cleared
3.2.5 Connection Time-out (Minutes)	30	30	30	30	30	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → IMAP4 → [Specific IMAP4 Virtual Server] → Properties → General Tab → Connection time-out (minutes)	Connection timeout limit (after which an inactive client is disconnected).	30
3.3 Delivery Restrictions									
3.3.1 Exclude or Limit Connections	Select "Only the list below" radio group. Leave the list blank.	Select "Only the list below" radio group. Leave the list blank.	Select "Only the list below" radio group. Enter authorized hosts.	Select "Only the list below" radio group. Leave the list blank.	Select "Only the list below" radio group. Leave the list blank.	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → IMAP4 → [Specific IMAP4 Virtual Server] → Properties → Access Tab → Connection control → Connection Button (Multiple Items)	Exclude or limit connections to a list of hosts.	None
4. Mailbox Store									
4.1 Authentication									
4.1.1 Clients Support S/MIME Signatures	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → General Tab → Clients support S/MIME signatures	Use this option to select the checkbox if the e-mail clients using this mailbox store support S/MIME.	Selected

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
4.2 Backup/Restore									
4.2.1 Time of Last Full Backup	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Database Tab → Time of last full backup	Provides the time when the last full backup occurred.	This type of backup was never performed.
4.2.2 Time of Last Incremental Backup	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Database Tab → Time of last incremental backup	Provides the time when the last incremental backup occurred.	This type of backup was never performed.
4.2.3 Do Not Permanently Delete Mailboxes Until Backed Up	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Limits Tab → Deletion settings → Do not permanently delete mailboxes and items until the store has been backed up	Do not permanently delete mailboxes and items until the store has been backed up.	Cleared
4.2.4 Maintenance Interval	Run daily for at least 4 hours	Run daily for at least 4 hours	Run daily for at least 4 hours	Run daily for at least 4 hours	Run daily for at least 4 hours	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Database Tab → Maintenance interval	Maintenance Interval	Run daily from 1:00am to 5:00am.
4.2.5 This Database can be Overwritten by a Restore	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Database Tab → This database can be overwritten by a restore	This database can be overwritten by a restore.	Cleared
4.3 Deletion Settings									
4.3.1 Do Not Permanently Delete Mailboxes Until Backed Up							See item 4.2.3	See item 4.2.3	

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
4.3.2 Keep Deleted Items for (days)	7	7	7	7	7	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Limits Tab → Deletion settings → Keep deleted items for (days)	Delete settings: keep deleted items for [days].	7
4.3.3 Keep Deleted Mailboxes for (days)	30	30	30	30	30	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Limits Tab → Deletion settings → Keep deleted mailboxes for (days)	Keep deleted mailboxes for [days].	30
4.4 Display									
4.4.1 Mailbox Store Policies	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Policies Tab	Used to list the policies that have been applied to the mailbox store.	Empty
4.5 Logging									
4.5.1 Archive All Messages Sent or Received by Mailboxes	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → General Tab → Archive all messages sent or received by mailboxes on this store	Archive all messages sent or received by mailboxes on this store.	Cleared
4.6 Permissions									
4.6.1 Mailbox Store Permissions	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Security Tab	Control permissions to the store.	Lists all users and their permissions; can add/remove users.
4.7 Storage Limits									
4.7.1 Storage Limits	Select all checkboxes	Select all checkboxes	Select all checkboxes	Select all checkboxes	Select all checkboxes	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Limits Tab → Storage limits (Multiple Items)	Storage limits: issue warning at [kb] ; prohibit send at [kb] ; prohibit send and receive at [kb]	All cleared.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
4.8 Mounting Stores									
4.8.1 Do Not Mount This Store at Start-up	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Mailbox Store [server] → Properties → Database Tab → Do not mount this store at start-up	Do not mount this store at start-up.	Cleared
5. NNTP									
5.1 Admin Post Requests									
5.1.1 Allow Control Messages	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → Settings Tab → Allow control messages	Determine whether control messages can be used to perform simple administrative functions without direct oversight.	Cleared
5.2 Authentication									
5.2.1 Authentication Method Used to Access NNTP Virtual Server	Select Basic authentication, Integrated Windows Authentication, Enable SSL client authentication, Require SSL client auth, ad Enable Client certificate mapping checkboxes.	Select Basic authentication, Integrated Windows Authentication, Enable SSL client authentication, Require SSL client auth, ad Enable Client certificate mapping checkboxes.	Select Basic authentication, Integrated Windows Authentication, Enable SSL client authentication, Require SSL client auth, ad Enable Client certificate mapping checkboxes.	Select Basic authentication, Integrated Windows Authentication, Enable SSL client authentication, Require SSL client auth, ad Enable Client certificate mapping checkboxes.	Select Basic authentication, Integrated Windows Authentication, Enable SSL client authentication, Require SSL client auth, ad Enable Client certificate mapping checkboxes.	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → Access Tab → Access control → Authentication Button (Multiple Items)	Determine the type of authentication to use and whether SSL should be required.	Basic authentication and Integrated Windows Authentication.
5.2.2 Certificate Wizard	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → Access Tab → Secure communication → Certificate Button → Wizard Button	Install or modify the virtual server certificate.	When clicked, the web server certificate wizard runs.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
5.3 Connections									
5.3.1 Require Secure Channel (SSL) and Require 128-bit encryption to NNTP Virtual Directory	Select both checkboxes	Select both checkboxes	Select both checkboxes	Select both checkboxes	Select both checkboxes	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Virtual Directories → [all entries here] → Properties → General Tab → Secure Communications → Secure Button → Require Secure Channel and Require 128-bit encryption	Control how the channel is secured.	When clicked, a window opens stating: Once a valid key certificate is installed, you can require that access take place on a secure channel.
5.3.2 Limit Number of Connections	5000	5000	5000	5000	5000	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → General Tab → Limit number of connections to	Limit number of connections to: (maximum number of simultaneous connections to the virtual server).	5000
5.3.3 Connection Time-out (Minutes)	10	10	10	10	10	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → General Tab → Connection time-out (minutes)	Connection time-out in minutes.	10
5.4 Delivery Restrictions									
5.4.1 Grant or Deny Access	Select "Only list below" radio button and then leave the list blank	Select "Only list below" radio button and then leave the list blank	Select "Only list below" radio button and then leave the list blank	Select "Only list below" radio button, and then specify a list of IPs which are permitted access to this NNTP server.	Select "Only list below" radio button and then leave the list blank	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → Access Tab → Connection control → Connection Button	Grant or deny access to this resource using IP addresses or Internet domain names.	
5.4.2 Feed: Limit Post Size	Select: 1500KB (if feed posting is enabled)	Select: 1500KB (if feed posting is enabled)	Select: 1500KB (if feed posting is enabled)	Select: 1500KB (if feed posting is enabled)	Select: 1500KB (if feed posting is enabled)	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → Settings Tab → Limit post size (KB)	The maximum size of articles that can be retrieved via feed posting.	Selected - 1500 KB.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
5.4.3 Feed: Limit Connection Size	Select: 40 MB (if feed posting is enabled)	Select: 40 MB (if feed posting is enabled)	Select: 40 MB (if feed posting is enabled)	Select: 40 MB (if feed posting is enabled)	Select: 40 MB (if feed posting is enabled)	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → Settings Tab → Limit connection size (MB)	The maximum amount of data that can be transferred in a single session via feed posting.	Selected - 40 MB.
5.4.4 Moderated	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → Newsgroups→ [all entries] → Properties → General Tab → Moderated	Specify whether some party should review all articles before they are posted to the newsgroup.	Cleared
5.4.5 Allow Posting	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Virtual Directories → [all entries here] → Properties → General Tab → Access restrictions → Allow posting	NNTP access restrictions.	Allow posting selected.
5.4.6 Restrict Newsgroup Visibility	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Virtual Directories → [all entries here] → Properties → General Tab → Access restrictions → Restrict newsgroup_visibility	NNTP access restrictions.	Restrict newsgroup_visibility cleared.
5.4.7 Allow Client Posting: Limit Post Size (kb)	100K	100K	100K	100K	100K	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → Settings Tab → Allow client posting	Allow client posting: limit post size (kb)	1000
5.4.8 Allow Client Posting: Limit Connection Size (mb)	1MB	1MB	1MB	1MB	1MB	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → Settings Tab → Allow client posting → Limit connection size (KB)	Allow client posting: limit connection size (mb)	20

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
5.5 Directory Access									
5.5.1 Physical Location Security Panel	Restrict all privileges to Administrator and Exchange Admin	Restrict all privileges to Administrator and Exchange Admin	Restrict all privileges to Administrator and Exchange Admin	Restrict all privileges to Administrator and Exchange Admin (if Control Newsgroup is to be used, add write privileges to those who may post to it.)	Restrict all privileges to Administrator and Exchange Admin	Manual	Windows Properties - NNTP Control Directory Physical Location Security Panel	Exchange can restrict NNTP directory access on a per user basis.	Multiple values (permissions window).
5.6 Logging									
5.6.1 Enable Logging	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Properties → General Tab → Enable logging	Set whether activity should be logged.	Cleared
5.6.2 Log Access	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → NNTP → [specific NNTP Virtual Server] → Virtual Directories → [all entries here] → Properties → General Tab → Content control → Log access	Determine whether activity should be logged.	Log access and index news content both selected.
6. POP3									
6.1 Authentication									
6.1.1 Authentication Method used to Access POP3 Virtual Directory	Select the "Basic authentication" and "SSL/TLS" checkboxes	Select the "Basic authentication" and "SSL/TLS" checkboxes	Select the "Basic authentication" and "SSL/TLS" checkboxes	Select the "Basic authentication" and "SSL/TLS" checkboxes	Select the "Basic authentication" and "SSL/TLS" checkboxes	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → POP3 → [Specific POP3 Virtual Server] → Properties → Access Tab → Access control → Authentication Button	Authentication (which authentication methods should be enabled).	Basic and simple (NTLM).

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
6.1.2 Certificate Wizard	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → POP3 → [Specific POP3 Virtual Server] → Properties → Access Tab → Secure communication → Certificate Button → Wizard Button	Opens a wizard to create and install certificate.	When clicked, the web server certificate wizard runs.
6.2 Connections									
6.2.1 Require Secure Channel (SSL) and Require 128-bit encryption to POP3 Virtual Server	Select both checkboxes	Select both checkboxes	Select both checkboxes	Select both checkboxes	Select both checkboxes	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → POP3 → [Specific POP3 Virtual Server] → Properties → Access Tab → Secure communication → Communication Button → Require Secure Channel and Require 128-bit encryption	Require a secure channel (between POP servers).	Not available until certificate is added.
6.2.2 Use SSL Connections	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → POP3 → [Specific POP3 Virtual Server] → Properties → Calendaring Tab → Use SSL connections	Use of SSL when downloading meeting requests.	Cleared
6.2.3 TCP Port/SSL Port	Enter 110/995 respectively .	Enter 110/995 respectively .	Enter 110/995 respectively .	Enter 110/995 respectively .	Enter 110/995 respectively .	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → POP3 → [Specific POP3 Virtual Server] → Properties → General Tab → Advanced Button → Edit Button → TCP port and SSL port	Specify which ports are used for normal and SSL connections.	110/995 respectively.
6.2.4 Limit Number of Connections	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → POP3 → [Specific POP3 Virtual Server] → Properties → General Tab → Limit number of connections to	Limit number of connections to: (maximum number of simultaneous connections to the virtual server).	Cleared
6.2.5 Connection Time-out (Minutes)	10	10	10	10	10	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → POP3 → [Specific POP3 Virtual Server] → Properties → General Tab → Connection time-out (minutes)	Connection timeout limit (after which an inactive client is disconnected).	10

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
6.3 Delivery Restrictions									
6.3.1 Exclude or Limit Connections	Select "Only the list below" radio group. Leave the list blank.	Select "Only the list below" radio group. Enter authorized hosts.	Select "Only the list below" radio group. Leave the list blank.	Select "Only the list below" radio group. Leave the list blank.	Select "Only the list below" radio group. Leave the list blank.	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → POP3 → [Specific POP3 Virtual Server] → Properties → Access Tab → Connection control → Connection Button	Exclude or limit connections to a list of hosts.	None
7. Public Folder Store									
7.1 Backup/Restore									
7.1.1 Time of Last Full Backup	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Database Tab → Time of last full backup	Provides the time when the last full backup occurred.	This type of backup was never performed.
7.1.2 Time of Last Incremental Backup	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Database Tab → Time of last incremental backup	Provides the time when the last incremental backup occurred.	This type of backup was never performed.
7.1.3 Maintenance Interval	Run daily for at least 4 hours	Run daily for at least 4 hours	Run daily for at least 4 hours	Run daily for at least 4 hours	Run daily for at least 4 hours	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store → Properties → Database Tab → Maintenance interval	Maintenance interval	Run daily from 1:00am to 5:00am.
7.1.4 This Database Can be Overwritten by a Restore	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Database Tab → This database can be overwritten by a restore	This database can be overwritten by a restore.	Cleared
7.1.5 Do Not Permanently Delete Items until Backed-Up	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Limits Tab → Deletion settings → Do not permanently delete items until the store has been backed up	Deletion settings: do not permanently delete items until the store has been backed up.	Cleared

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
7.2 Connections									
7.2.1 Clients Support S/MIME Signatures	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → General Tab → Clients support S/MIME signatures	Use this option to select the checkbox if the e-mail clients using this mailbox store support S/MIME.	Selected
7.3 Deletion Settings									
7.3.1 Keep Deleted Items for (days)	7	7	7	7	7	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Limits Tab → Deletion settings → Keep deleted items for (days)	Deletion settings: keep deleted items for [days].	7
7.3.2 Do Not Permanently Delete Items until Backed-Up							See item 7.1.5	See item 7.1.5.	
7.3.3 Age Limit for All Folders in This Store (Days)	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Limits Tab → Age limits → Age limit for all folders in this store (days)	Age limits: age limit for all folders in this store {days}.	Cleared
7.4 Display									
7.4.1 Public Folder Policies	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Policies Tab	Used to list the policies that have been applied to the folder store.	Blank
7.5 Mounting Stores									
7.5.1 Do Not Mount This Store at Start-up	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Database Tab → Do not mount this store at start-up	Do not mount this store at start-up.	Cleared

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
7.6 Permissions									
7.6.1 Public Folder Permissions	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Security Tab	Group or user names and permissions.	Lists all users and their permissions; can add/remove users.
7.7 Storage Limits									
7.7.1 Issue Warning at (kb)/Prohibit Post at (kb)	Enter values for both "Issue Warning at" and "Prohibit Post at" (Specific values will be enterprise dependent.)	Enter values for both "Issue Warning at" and "Prohibit Post at" (Specific values will be enterprise dependent.)	Enter values for both "Issue Warning at" and "Prohibit Post at" (Specific values will be enterprise dependent.)	Enter values for both "Issue Warning at" and "Prohibit Post at" (Specific values will be enterprise dependent.)	Enter values for both "Issue Warning at" and "Prohibit Post at" (Specific values will be enterprise dependent.)	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Limits Tab → Storage limits → Issue warning at (KB) and Prohibit post at (KB)	Storage limits: issue warning at [kb] and prohibit post at [kb].	Blank
7.7.2 Maximum Item Size (kb)	Select the checkbox and apply maximum item size of 10240	Select the checkbox and apply maximum item size of 10240	Select the checkbox and apply maximum item size of 10240	Select the checkbox and apply maximum item size of 10240	Select the checkbox and apply maximum item size of 10240	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Public Folder Store [server] → Properties → Limits Tab → Maximum item size (KB)	Storage limits: Maximum item size [kb].	Selected and size 10240.
8. Public Folders									
8.1 Deletion Settings									
8.1.1 Deletion Settings	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Limits Tab → Deletion settings → Use public store defaults	Deletion settings (use public store defaults/ specify settings).	Selected (use default).
8.1.2 Age Limits	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Limits Tab → Age limits → Use public store defaults	Age limits (use public store defaults or set).	Selected (use default).

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
8.2 Delivery Restrictions									
8.2.1 Sending Message Size	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Exchange General Tab → Delivery Restrictions Button → Sending messages size → Use default limit	Restrict sending message size from this folder.	Use Default Limit (blank).
8.2.2 Receiving Message Size	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Exchange General Tab → Delivery Restrictions Button → Receiving message size → Use default limit	Restrict receiving message size to this folder.	Use Default Limit (blank).
8.2.3 Message Restrictions	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Exchange General Tab → Delivery Restrictions Button → Message restrictions → Accept messages	Restrict who may send messages to this folder.	Accept Messages from Everyone.
8.3 Permissions									
8.3.1 Client Permissions	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Permissions Tab → Client permissions Button	Permissions assigned to roles and Exchange recipients assigned to roles.	See description.
8.3.2 Directory Rights	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Permissions Tab → Directory rights Button	Permissions assigned to roles and Exchange recipients assigned to roles.	See description.
8.3.3 Administrative Rights	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Permissions Tab → Administrative rights Button	Permissions assigned to roles and Exchange recipients assigned to roles.	See description.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
8.4 Storage Limits									
8.4.1 Storage Limits	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	Select the "Use public store defaults" checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Limits Tab → Storage limits (Multiple Items)	Storage limits (use public store defaults and specify settings).	All cleared.
8.5 Visibility									
8.5.1 Address List Name	Select radio button "Use this name" and fill in a name for the folder distinct from the folder name.	Select radio button "Use this name" and fill in a name for the folder distinct from the folder name.	Select radio button "Use this name" and fill in a name for the folder distinct from the folder name.	Select radio button "Use this name" and fill in a name for the folder distinct from the folder name.	Select radio button "Use this name" and fill in a name for the folder distinct from the folder name.	S	Exchange System Manager → Administrative Group → [administrative group] → Folders → Public Folders → [public folder] → Properties → General Tab → Address list name	Change the display name of the public folder. Options are to use existing folder name or create an alias for the folder.	Use folder name.
8.5.2 Hide from Exchange Address List	Select the "Hide from Exchange address lists" checkbox	Select the "Hide from Exchange address lists" checkbox	Select the "Hide from Exchange address lists" checkbox	Select the "Hide from Exchange address lists" checkbox	Select the "Hide from Exchange address lists" checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Exchange Advanced Tab → Hide from Exchange address lists	Hide from Exchange address lists (do not display this address in lists).	Selected
8.5.3 Send on Behalf	None	None	None	None	None	S	Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folders → [public folder] → Properties → Exchange General Tab → Delivery Options Button → Send on behalf	Send on behalf (which users can send mail on behalf of this folder).	None
9. Routing Groups									
9.1 Authentication									
9.1.1 Authentication and Encryption Algorithms	Select the "Basic authentication " and "TLS encryption" checkboxes	Select the "Basic authentication " and "TLS encryption" checkboxes	Select the "Basic authentication " and "TLS encryption" checkboxes	Select the "Basic authentication " and "TLS encryption" checkboxes	Select the "Basic authentication " and "TLS encryption" checkboxes	S	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [SMTP connector] → Properties → Advanced Tab → Outbound Security Button	Access security (Anonymous, Basic, Integrated Windows Authentication).	Anonymous

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
9.2 Connections									
9.2.1 Connector Scope	Select the "Routing group" checkbox	Select the "Routing group" checkbox	Select the "Routing group" checkbox	Select the "Routing group" checkbox	Select the "Routing group" checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [SMTP connector] → Properties → Address Space Tab → Connector Scope	Connector scope (entire organization/routing group).	Entire Organization.
9.2.2 Routing Options	Specify the name of the smart host (if configured)	Specify the name of the smart host (if configured)	Specify the name of the smart host (if configured)	Specify the name of the smart host (if configured)	Specify the name of the smart host (if configured)	R	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [SMTP connector] → Properties → General Tab → Radio Group	Specify whether to send the mail to a specific host which will handle routing (a "smart host"), or whether this virtual server should contact a DNS server and handle routing itself.	Use DNS.
9.3 Delivery Restrictions									
9.3.1 Accept or Reject Outbound Messages through this Routing Group Connector	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [routing group connector] → Properties → Delivery Restrictions	Accepted/Rejected (default action and a list of exceptions).	Accepted (no exception).
9.3.2 Servers Allowed to Send Mail over Specific SMTP Connector	Select "These servers can send mail over this connector" checkbox and list allowed SMTPVSs	Select "These servers can send mail over this connector" checkbox and list allowed SMTPVSs	Select "These servers can send mail over this connector" checkbox and list allowed SMTPVSs	Select "These servers can send mail over this connector" checkbox and list allowed SMTPVSs	Select "These servers can send mail over this connector" checkbox and list allowed SMTPVSs	S	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [routing group connector] → Properties → General Tab → Any local server can send mail over this connector/These servers can send mail over this connector	These servers can send mail over this connector/any local servers can send mail over this connector.	Any local server.
9.3.3 Allow Messages to be Relayed over this SMTP connector to These Domains	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [SMTP connector] → Properties → Address Space Tab → Allow messages to be relayed to these domains	Allow messages to be relayed to these domains.	Cleared

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
9.3.4 Accept or Reject Outbound Messages through this SMTP Connector	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [SMTP connector] → Properties → Delivery Restrictions Tab	Accepted/Rejected (default action and a list of exceptions).	Accepted (no exception).
9.3.5 Allowed Sizes for Routing Group Connector Messages	Clear the checkbox. No limit	Clear the checkbox. No limit	Clear the checkbox. No limit	Clear the checkbox. No limit	Clear the checkbox. No limit	S	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [routing group connector] → Properties → Content Restrictions Tab → Allowed sizes → Only messages less than (KB)	Allowed sizes.	No limit.
9.3.6 Allowed Types for SMTP Connector Messages	If SMTP connectors are used, use multiple connectors with at least one connector dedicated exclusively to system messages. Be sure, however, that both system and non-system messages can travel between all SMTP servers	If SMTP connectors are used, use multiple connectors with at least one connector dedicated exclusively to system messages. Be sure, however, that both system and non-system messages can travel between all SMTP servers	If SMTP connectors are used, use multiple connectors with at least one connector dedicated exclusively to system messages. Be sure, however, that both system and non-system messages can travel between all SMTP servers	If SMTP connectors are used, use multiple connectors with at least one connector dedicated exclusively to system messages. Be sure, however, that both system and non-system messages can travel between all SMTP servers	If SMTP connectors are used, use multiple connectors with at least one connector dedicated exclusively to system messages. Be sure, however, that both system and non-system messages can travel between all SMTP servers	R	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [SMTP connector] → Properties → Content Restrictions Tab → Allowed types (Multiple Items)	Allowed types (system [Exchange or Win03] vs. non-system [users, distribution groups, contacts]).	All types allowed.
9.3.7 Allowed Sizes for SMTP Connector Messages	Clear the checkbox. No limit	Clear the checkbox. No limit	Clear the checkbox. No limit	Clear the checkbox. No limit	Clear the checkbox. No limit	S	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [SMTP connector] → Properties → Content Restrictions Tab → Allowed sizes → Only messages less than (KB)	Allowed sizes.	No limit.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
9.3.8 Allowed Types for Routing Group Connector Messages	Use multiple connectors with at least one connector dedicated exclusively to system messages. Be sure, however, that both system and non-system messages can travel between routing groups.	Use multiple connectors with at least one connector dedicated exclusively to system messages. Be sure, however, that both system and non-system messages can travel between routing groups.	Use multiple connectors with at least one connector dedicated exclusively to system messages. Be sure, however, that both system and non-system messages can travel between routing groups.	Use multiple connectors with at least one connector dedicated exclusively to system messages. Be sure, however, that both system and non-system messages can travel between routing groups.	Use multiple connectors with at least one connector dedicated exclusively to system messages. Be sure, however, that both system and non-system messages can travel between routing groups.	R	Exchange System Manager → Administrative Groups → [administrative group] → Routing Groups → [routing group] → Connectors → [routing group connector] → Properties → Content Restrictions Tab → Allowed types	Allowed types (system [Exchange or Win03] vs. non-system [users, distribution groups, contacts]).	All types allowed.
10. Servers									
10.1 Deletion Settings									
10.1.1 Zero Out Deleted Database Pages	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Properties → General Tab → Zero out deleted database pages	More secure deletion of databases by overwriting the memory in which they were stored.	Cleared
10.2 Display									
10.2.1 Server Policy Display	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Policies Tab	Display the policies on the server.	N/A
10.3 Logging									
10.3.1 Enable Circular Logging	Clear the checkbox on back-end server, select the checkbox on front-end server	Clear the checkbox on back-end server, select the checkbox on front-end server	Clear the checkbox on back-end server, select the checkbox on front-end server	Clear the checkbox on back-end server, select the checkbox on front-end server	Clear the checkbox on back-end server, select the checkbox on front-end server	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → [storage group] → Properties → General Tab → Enable circular logging	Determine whether circular logging is enabled (overwriting of old log files as opposed to the continuous creation of new files).	Cleared

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
10.3.2 Diagnostic Log Levels	Select "None"	Select "None"	Select "None"	Select "None"	Select "None"	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Diagnostic Logging → Logging level	Set diagnostic log levels.	No audits enabled.
10.3.3 Enable Subject Logging and Display	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → General Tab → Enable subject logging and display	When message tracking is enabled, include the message subject.	Cleared
10.3.4 Enable Message Tracking	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → General Tab → Enable message tracking	Enable message tracking. Log messages until they are handed off to SMTPVS or after receiving messages from SMTPVS.	Cleared
10.3.5 Remove Log Files (days)	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → General Tab → Remove log files	Determine how many days logs should be kept before deletion.	Cleared
10.4 Monitoring and Tools									
10.4.1 Monitoring Information	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Queues	Info and controls for the various message delivery queues.	N/A

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
10.4.2 Display Domain Controllers	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Directory Access Tab	Displays the domain controllers (DC, Global Catalog, and configuration servers - no controls).	All Domain Controllers shown.
10.4.3 Default Microsoft Exchange Services	NR	NR	NR	NR	NR	Manual	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Monitoring Tab → Default Microsoft Exchange Services → Detail Button	Displays the services & their status (can add or remove, but not control services from here).	N/A
10.4.4 Automatically Send Fatal Service Error Information to Microsoft	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → General Tab → Automatically send fatal service error information to Microsoft	Automatically send messages to Microsoft when system errors are detected.	Cleared
10.4.5 Disable all Monitoring of This Server	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Monitoring Tab → Disable all monitoring of this server	Disable all monitoring of this server.	Cleared
10.4.6 Warning State/Critical State of Disk Space Threshold	Warning 3GB Critical 2GB	Warning 3GB Critical 2GB	Warning 3GB Critical 2GB	Warning 3GB Critical 2GB	Warning 3GB Critical 2GB	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Monitoring Tab → Disk Space Threshold → Details Button → Warning State and Critical State	Determine notification triggers for disk space events.	Not configured.
10.4.7 Duration / Warning State / Critical State of CPU Utilization	Duration 10 Warning 80 Critical 90	Duration 10 Warning 80 Critical 90	Duration 10 Warning 80 Critical 90	Duration 10 Warning 80 Critical 90	Duration 10 Warning 80 Critical 90	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Monitoring Tab → CPU Utilization Threshold → Details Button → Duration and Warning State and Critical State	Determine notification triggers for CPU threshold events.	Not configured.
10.4.8 Warning State / Critical State of SMTP Queues	Warning 10 Critical 20	Warning 10 Critical 20	Warning 10 Critical 20	Warning 10 Critical 20	Warning 10 Critical 20	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Monitoring Tab → SMTP Queue Threshold → Details Button → Warning State and Critical State	Determine notification triggers for SMTP queue growth events.	Not configured.

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
10.4.9 Warning State / Critical State of X.400 Queue Thresholds	Warning 10 Critical 20	Warning 10 Critical 20	Warning 10 Critical 20	Warning 10 Critical 20	Warning 10 Critical 20	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Monitoring Tab → X.400 Queue Thresholds → Details Button → Warning State Critical State	Determine notification triggers for X.400 queue growth events.	Not configured.
10.4.10 Duration / Warning State / Critical State of Virtual Memory Threshold	Duration 3 Warning 25 Critical 10	Duration 3 Warning 25 Critical 10	Duration 3 Warning 25 Critical 10	Duration 3 Warning 25 Critical 10	Duration 3 Warning 25 Critical 10	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Monitoring Tab → Virtual Memory Threshold → Details Button → Duration and Warning State and Critical State	Determine notification triggers for virtual memory events.	Not configured.
10.4.11 When Windows 2000 Service is Not Running, Change State to	Critical; add Event Log, NT LM Security Support Provider, Remote Procedure Call (RPC), Server, Workstation, IIS Admin Services, HTTP SSL	Critical; add Event Log, NT LM Security Support Provider, Remote Procedure Call (RPC), Server, Workstation, IIS Admin Services, Microsoft Exchange POP3	Critical; add Event Log, NT LM Security Support Provider, Remote Procedure Call (RPC), Server, Workstation, IIS Admin Services, Microsoft Exchange IMAP4	Critical; add Event Log, NT LM Security Support Provider, Remote Procedure Call (RPC), Server, Workstation, IIS Admin Services, Network News Transfer Protocol (NNTP)	Critical; add Event Log, NT LM Security Support Provider, Remote Procedure Call (RPC), Server, Workstation, IIS Admin Services	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Monitoring Tab → [windows 2000 Service] → Details Button → When service is not running change state to	When service stops running change state to: (Admin can select whether “critical” or “warning” should be displayed when a service is not running).	Not configured.
10.4.12 When Core Exchange Services are Not Running, Change State to	Critical	Critical (Remove World Wide Web Publishing Service)	Critical (Remove World Wide Web Publishing Service)	Critical (Remove World Wide Web Publishing Service)	Critical (Remove World Wide Web Publishing Service)	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Properties → Monitoring Tab → Default Microsoft Exchange Services→ Details Button → When service is not running change state to	When service stops running change state to: (Admin can select whether “critical” or “warning” should be displayed when a service is not running).	Critical

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
10.4.13 Notification Creation	Create notifications for warning and critical triggers for all servers and notifications for the down trigger for all connectors.	Create notifications for warning and critical triggers for all servers and notifications for the down trigger for all connectors.	Create notifications for warning and critical triggers for all servers and notifications for the down trigger for all connectors.	Create notifications for warning and critical triggers for all servers and notifications for the down trigger for all connectors.	Create notifications for warning and critical triggers for all servers and notifications for the down trigger for all connectors.	S	Exchange System Manager → Tools → Monitoring and Status → Notifications	Create notifications that go out when the monitor detects warning or critical events.	None created.
11. SMTP									
11.1 Authentication									
11.1.1 Authentication Method Used for Access to SMTP Virtual Directory	Select only "Basic authentication" and "TLS encryption"	Select only "Basic authentication" and "TLS encryption"	Select only "Basic authentication" and "TLS encryption"	Select only "Basic authentication" and "TLS encryption"	Select only "Basic authentication" and "TLS encryption"	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Access Tab → Access control → Authentication Button	Admin selection of type of permitted authentication (Anonymous, Basic (TLS), Integrated Windows Authentication).	All enabled (except TLS).
11.1.2 Resolve Anonymous Email	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Access Tab → Access Control → Authentication Button → Resolve anonymous email	Resolve anonymous email (reverse DNS lookups on mail domains).	Cleared
11.1.3 Certificate Wizard	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	Execute Wizard to Install Certificate	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Access Tab → Secure communication → Certificate Button → Wizard Button	Opens a wizard to create and install certificate.	When clicked, the web server certificate wizard runs.
11.1.4 Authentication Method Used for Outbound Connections	Select "Anonymous" and "TLS encryption"	Select "Anonymous" and "TLS encryption"	Select "Anonymous" and "TLS encryption"	Select "Anonymous" and "TLS encryption"	Select "Anonymous" and "TLS encryption"	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Delivery Tab → Outbound Security Button	Authentication method for outbound connections.	Anonymous

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
11.2 Connections									
11.2.1 Require Secure Channel (SSL) and Require 128 bit encryption to SMTP Virtual Server	Select both checkboxes	Select both checkboxes	Select both checkboxes	Select both checkboxes	Select both checkboxes	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Access Tab → Secure communication → Communication Button → Require secure channel and Require 128-bit encryption	Require a secure channel (between SMTP servers).	Not available until certificate is added.
11.2.2 Exclude or Limit Relaying SMTP Servers	Select "Only the list below"	Select "Only the list below"	Select "Only the list below"	Select "Only the list below"	Select "Only the list below"	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Access Tab → Relay restrictions → Relay Button	Exclude or limit relaying SMTP servers to a list of hosts.	"Only list below" selected.
11.2.3 Exclude or Limit Relaying SMTP Servers	Select "Allow all computers which successfully authenticate ..."	Select "Allow all computers which successfully authenticate ..."	Select "Allow all computers which successfully authenticate ..."	Select "Allow all computers which successfully authenticate ..."	Select "Allow all computers which successfully authenticate ..."	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Access Tab → Relay restrictions → Relay Button	Exclude or limit relaying SMTP servers to a list of hosts.	None
11.2.4 Smart Host	None	None	None	None	None	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Delivery Tab → Advanced Button → Smart host	Specify the smart host that should be used to relay outgoing messages.	None
11.2.5 Perform Reverse DNS Lookup on Incoming Messages	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Delivery Tab → Advanced Button → Perform reverse DNS lookup on incoming messages	Perform reverse DNS lookup on incoming messages.	Cleared
11.2.6 TCP Port for Outbound Connections	25	25	25	25	25	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Delivery Tab → Outbound connections Button → TCP port	Specify a TCP port (for outbound connections).	25

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
11.2.7 TCP Port for SMTP Virtual Server	25	25	25	25	25	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → General Tab → Advanced → Edit → TCP port	Changing IP port for SMTPVS.	25
11.2.8 Outbound Delivery Retry (group)	First retry: 10 minutes, Second retry: 15 minutes, Third retry: 15 minutes, Subsequent retry: 15 minutes, Delay notification: 12 hours, Expiration timeout: 2 days	First retry: 10 minutes, Second retry: 15 minutes, Third retry: 15 minutes, Subsequent retry: 15 minutes, Delay notification: 12 hours, Expiration timeout: 2 days	First retry: 10 minutes, Second retry: 15 minutes, Third retry: 15 minutes, Subsequent retry: 15 minutes, Delay notification: 12 hours, Expiration timeout: 2 days	First retry: 10 minutes, Second retry: 15 minutes, Third retry: 15 minutes, Subsequent retry: 15 minutes, Delay notification: 12 hours, Expiration timeout: 2 days	First retry: 10 minutes, Second retry: 15 minutes, Third retry: 15 minutes, Subsequent retry: 15 minutes, Delay notification: 12 hours, Expiration timeout: 2 days	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Delivery Tab → Outbound (Multiple Items)	Outbound Delivery retry intervals & timeouts (in case of mail that has trouble being sent).	First retry: 10 minutes, Second retry: 15 minutes, Third retry: 15 minutes, Subsequent retry: 15 minutes, Delay notification: 12 hours, Expiration timeout: 2 days.
11.2.9 Maximum Hop Count	30	30	30	30	30	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Delivery Tab → Advanced Button → Maximum hop count	Maximum hop count.	30
11.2.10 Limit Number of Connections	Select the checkbox and apply 1000 as the limit.	Select the checkbox and apply 1000 as the limit.	Select the checkbox and apply 1000 as the limit.	Select the checkbox and apply 1000 as the limit.	Select the checkbox and apply 1000 as the limit.	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Delivery Tab → Outbound connections Button → Limit number of connections to	Limit number of connections (limit total outbound connections).	Selected (1000)
11.2.11 Time Out (minutes)	10	10	10	10	10	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Delivery Tab → Outbound connections Button → Time-out (minutes)	Time out (for outbound connections).	Selected (10)

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
11.2.12 Limit Number of Connections Per Domain	Select the checkbox and apply 100 as the limit.	Select the checkbox and apply 100 as the limit.	Select the checkbox and apply 100 as the limit.	Select the checkbox and apply 100 as the limit.	Select the checkbox and apply 100 as the limit.	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Delivery Tab → Outbound connections Button → Limit number of connections per domain	Limit number of connections per domain.	Selected (100)
11.2.13 Limit Number of Simultaneous Connections to Virtual Server	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → General Tab → Limit number of connections to	Maximum number of simultaneous connections to the virtual server.	Cleared (no limit).
11.2.14 Connection Time-out (minutes)	10 minutes	10 minutes	10 minutes	10 minutes	10 minutes	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → General Tab → Connection time-out (minutes)	Connection timeout limit (after which an inactive client is disconnected).	10 minutes.
11.3 Delivery Restrictions									
11.3.1 Exclude or Limit Connections	Select "Only the list below" and enter a list of trusted servers.	Select "Only the list below" and enter a list of trusted servers	Select "Only the list below" and enter a list of trusted servers	Select "Only the list below" and enter a list of trusted servers	Select "Only the list below" and enter a list of trusted servers	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Access Tab → Connection control → Connection Button	Exclude or limit connections to a list of hosts.	None
11.3.2 Apply Sender/Recipient/Connection Filters	Select all checkboxes	Select all checkboxes	Select all checkboxes	Select all checkboxes	Select all checkboxes	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → General Tab → Advanced → Edit → Apply Sender Filter/Apply Recipient Filter/Apply Connection Filter	Sender and Recipient and Connection Filter Buttons.	Cleared
11.3.3 Limit Message Size	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Messages Tab → Limit message size to: (KB)	Limit message size (max size of messages to or from the server, including attachments).	Cleared

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
11.3.4 Limit Session Size	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	Clear the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Messages Tab → Limit session size to: (KB)	Limit session size (max size of data to or from the server in a single session).	Cleared
11.3.5 Limit Number of Messages Per Connection	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Messages Tab → Limit number of messages per connection to	Limit number of messages per connection (messages to the virtual server from other SMTP hosts).	Selected (20)
11.3.6 Limit Number of Recipients per Message	NR	NR	NR	NR	NR	R	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → Messages Tab → Limit number of recipients per message to	Limit number of recipients per message (for any message through the virtual server).	Selected (64000)
11.4 Logging									
11.4.1 Enable Logging	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	Select the checkbox	S	Exchange System Manager → Administrative Groups → [administrative group] → Servers → [server] → Protocols → SMTP → [Specific SMTP Virtual Server] → Properties → General Tab → Enable logging	Enable logging of connections between SMTP hosts, and also specify the log format.	Disabled
12. User Level Controls									
12.1 Deletion Settings									
12.1.1 Deleted Item Retention (group)	Select the "Use Mailbox Store Defaults" checkbox	Select the "Use Mailbox Store Defaults" checkbox	Select the "Use Mailbox Store Defaults" checkbox	Select the "Use Mailbox Store Defaults" checkbox	Select the "Use Mailbox Store Defaults" checkbox	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange General Tab → Storage Limits Button → Deleted item retention → Use mailbox store defaults	Deleted item retention (use mail store defaults/specify settings).	Use Mailbox Store Defaults.
12.2 Delivery Restrictions									
12.2.1 Send on Behalf	None	None	None	None	None	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange General Tab → Delivery Options Button → Send on behalf	Send on behalf (which users can send mail on behalf of this mailbox).	None

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
12.2.2 Recipient Limits	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange General Tab → Delivery Options Button → Recipient limits	Recipient limits (max number of recipients for a single message from this user).	Use Default Limit (blank).
12.2.3 Sending Message Size/Receiving Message Size	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Select the "Use default limit" checkbox	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange General Tab → Delivery Restrictions Button → Sending message size and Receiving message size (Multiple Items)	Sending|receiving message size.	Use Default Limit (blank).
12.2.4 Message Restrictions (group)	Select the "From everyone" checkbox	Select the "From everyone" checkbox	Select the "From everyone" checkbox	Select the "From everyone" checkbox	Select the "From everyone" checkbox	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange General Tab → Delivery Restrictions Button → Message restrictions → Accept messages	Message restrictions (whether user is limited in who they can receive messages from).	Receive from everyone.
12.2.5 Accept Messages: From Authenticated Users Only	Select the "From authenticated users only" checkbox	Select the "From authenticated users only" checkbox	Select the "From authenticated users only" checkbox	Select the "From authenticated users only" checkbox	Select the "From authenticated users only" checkbox	Manual	Active Directory Users and Computers → [domain] → Users → [group] → Properties → Exchange General Tab → Delivery Restrictions Button → Message restrictions → Accept messages → From authenticated users only	Only allow this group to receive messages from authenticated (internal) sources.	Cleared
12.3 Enabled Services									
12.3.1 Enable POP3	NR	NR	NR	NR	NR	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange Features Tab	POP3 (specify whether the protocol is enabled and whether it uses defaults for this user).	Enabled using defaults.
12.3.2 Enable IMAP	NR	NR	NR	NR	NR	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange Features Tab	IMAP (specify whether the protocol is enabled and whether it uses defaults for this user).	Enabled using defaults.
12.3.3 Enable OWA	NR	NR	NR	NR	NR	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange Features Tab	Outlook Web access (specify whether protocol is enabled).	Enabled using defaults.
12.3.4 Enable OMA	NR	NR	NR	NR	NR	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange Features Tab	Outlook Mobile Access (specify whether service is enabled).	Enabled
12.3.5 Enable ActiveSync	NR	NR	NR	NR	NR	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange Features Tab	User initiated synchronization (whether ActiveSync is enabled for this user).	Enabled

Reference ID/Control Title	HTTP Server	POP3 Server	IMAP4 Server	NNTP Server	SMTP Server	S/R	GUI Location in Exchange	Description	Default
12.3.6 Up-to-Date Notifications	NR	NR	NR	NR	NR	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange Features Tab	Up-to-date Notifications (whether the server will send notices to user that ActiveSync should run).	Enabled
12.4 Permissions									
12.4.1 Mailbox Rights	Leave at default	Leave at default	Leave at default	Leave at default	Leave at default	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange Advanced Tab → Mailbox Rights Button	Rights to this mailbox.	SELF granted full access.
12.5 Storage Limits									
12.5.1 Storage Limits (groups)	Select the "Use Mailbox Store Defaults" checkbox	Select the "Use Mailbox Store Defaults" checkbox	Select the "Use Mailbox Store Defaults" checkbox	Select the "Use Mailbox Store Defaults" checkbox	Select the "Use Mailbox Store Defaults" checkbox	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange General Tab → Storage Limits Button → Storage limits → Use mailbox store defaults	Storage limits (use mail store defaults and specify settings).	Use Mailbox Store Defaults.
12.6 Visibility									
12.6.1 Hide from Exchange Address List	NR	NR	NR	NR	NR	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange Advanced Tab → Hide from Exchange address lists	Hide from Exchange address lists (do not display this address in lists).	Cleared
12.6.2 ILS Server Settings	None	None	None	None	None	Manual	Active Directory Users and Computers → [computer] → Users → [specific user] → Properties → Exchange Advanced Tab → ILS Settings Button	ILS Server & account for this user (to use Internet Locator Service).	None

Chapter

3

Rationales for Exchange Controls Security Settings

This chapter supplies rationales for the control-specific Exchange guidance found in Chapter 2. Rationales provide background material for a control's recommendation, a description of the control's functionality, the recommended value, and a synopsis of the expected ramifications if the recommended value is implemented or not implemented.

The Reference ID/Control Title column in Table 2 can be used as a primary key to obtain a one-to-one mapping between the rationales in Table 3 and the security recommendations in Table 2.

Table 3. Rationales – One to One Mapping to Security Recommendations

1. Global Settings
1.1 Automated Messages
1.1.1 Allow out-of-office Responses These items control the creation of automated messages. These include out-of-office responses, automated replies, mail forwarding, delivery reports, and non-delivery reports. Disable all of these features. In all cases (except mail forwarding), automated messages can be used by a third party to determine user liveness on the server. This can result in the disclosure of active user accounts to third parties, paving the way for possible future attacks. Out-of-office and automated replies are especially dangerous in this regard since they can alert a third party not only to the existence of a user's account but also that the user is not checking the account regularly and may not notice attempts to compromise this account. The mail forwarding features do not provide information to third parties, but pose a potential risk on networks where classified or confidential information may be sent. If auto-forwarding is configured, sensitive information sent to this user's account may automatically be transferred outside the control of the organization. For example, a sensitive memo sent to a distribution list that includes an account with auto-forward enabled, may be forwarded to an account that resides on a system that does not employ sufficient security to protect the message contents. While none of these features are a direct security risk, they may provide information to potential attackers and/or result in sensitive information being placed in an insecure environment. As a result, their use should be controlled. Note that disabling these features does mean that parties sending mail to users in the organization will not be informed by the organization if their mail could not be delivered. Note that this functionality can be configured based on the destination on the message. The "Default" format applies to all domains. However, if a new format is created and applied to a specific domain, that domain will use the new format's configuration while all other domains (those without specially designated formats) will use the Default format. In this way, it is possible to state that some domain (probably one's own domain) could be allowed to receive certain types of automatically generated messages. However, this is not recommended due to the risk of insider threats.

1.2 Delivery Restrictions

1.2.1 Global Accept and Deny List Configuration

These fields allow the administrator to list IP addresses and subnets from which messages should be accepted or denied, respectively. The settings in these fields override those specified in the Block List Service Configuration field (item 1.2.2). Moreover, the accept list overrides the deny list. (That is, an IP address that is present in both lists will always be accepted.) These controls are usually used to prevent connections from known sources of spam, but can also be used to block known malicious hosts.

List known sources of spam and other undesirable hosts in the deny list. Include your own domain, and the domains of important partners in the accept list. The latter should be done so that even if the block list to which you subscribe claims that one or more of these domains are spammers, you will still be able to receive mail from them.

Note that unless the “Apply connection filter” checkbox (item 11.3.2) is selected on an SMTP Virtual Server, this control will not take effect on that virtual server regardless of the setting of this control. Note also that these filters will only be applied to anonymous connections - users and computers that are able to authenticate will not be filtered using these rules.

1.2.2 Specify Block List Service Provider

This field is used to specify block list services. These services collect the IP addresses of known spammers and other hostile parties. Subscribers to these services can then use these addresses to filter out mail from these known sources of spam.

Subscribe to a block list service and configure Exchange to use it to filter messages. Block list servers have a fee. Sometimes the fee can be high so make sure the cost is affordable. Block list services take much of the work out of spam blocking since they update block lists automatically and maintain larger block lists than a single administrator could conveniently maintain. Failure to specify a block list will mean that administrators must manually specify each spammer address to block in the deny field (item 1.2.1) as it is discovered. One example of a block list provider is http://www.declude.com/Articles.asp?ID=97&Redirected=Y

Configure the block list service provider by clicking the “Add” button under the Block List Service Configuration list. Use the DNS Suffix of Provider field to specify the block list provider to which you have subscribed. (Your provider will provide a value for this field.)

Note that unless the “Apply connection filter” checkbox (item 11.3.2) is selected on a SMTP Virtual Server, this control will not take effect on that virtual server regardless of the settings of this control. Note also that these filters will only be applied to anonymous connections - users and computers that are able to authenticate will not be filtered using these rules.

1.2.3 Block List Exceptions

The Exception button is used to specify IP addresses that you do not want blocked despite their presence in a block list (item 1.2.2). This field can be used to specify the addresses of computers which have been incorrectly identified as spammers by the block list service and which you wish to receive mail from.

Carefully vet any addresses added to this field to ensure that they are, in fact, sources of legitimate mail.

Note that unless the “Apply connection filter” checkbox (item 11.3.2) is selected on the SMTP Virtual Server, this control will not take effect on that virtual server regardless of the setting of this control. Note also that these filters will only be applied to anonymous connections - users and computers that are able to authenticate will not be filtered using these rules.

1.2.4 Sending/Receiving Message Size

These fields control the maximum size of acceptable outbound and inbound messages, respectively. The outbound limit is to reduce network congestion and to minimize the chance of internal users sending inappropriately large messages to external parties. The inbound limit is to reduce the chance that external parties will be able to monopolize network bandwidth and disk space on the server through repeated sending of large messages.

While the precise limits of these settings can vary depending on the mission of the organization, message size limits should be set to at most 30 megabytes. This will accommodate most modern business requirements (for example, text documents, small-to-medium imagery files, briefing slides, and so on). Although it will not accommodate *all* business requirements (for example, video files, modeling data, and so on), it represents a starting point from which organizations can tune their policies. Note that messages larger than 30 megabytes will begin to slow down the functioning of the Exchange server if received frequently (although hardware and server usage considerations play heavily into this).

Some organizations may require one or both fields be given larger limits than those recommended above. If this occurs, extra vigilance on the part of the administrators may be needed to minimize abuse. However, some limit should be applied to both fields. Selecting the "no limit" radio button on either field is likely to result in abuse and can lead to rapid filling of server disk space. Note that these settings can be overridden on a per user/group basis using the Active Directory Users and Computer snap-in (item 12.2.3) This can be done if only specific users have legitimate need to send large email messages.

1.2.5 Recipient Limits

This field is used to control the maximum number of recipients that can be specified in a single message sent from this server. Its primary purpose is to minimize the chance of an internal sender spamming other recipients, since spam messages often have a large number of recipients.

While the precise value of this control may vary between organizations, the maximum number of recipients per message should be less than or equal to 5000. This value is probably larger than is needed for most organizations, but is small enough to minimize usefulness to spammers. Moreover, it represents a number of recipients that Exchange can easily handle.

A larger recipient limit may be specified if there is an organizational need to distribute messages with an extremely large mailing list (for example, if sending out customer alerts or for popular mailing lists). However, some limit on the number of recipients should be applied. Selecting the "no limit" radio button for this item is likely to result in abuse.

Note that these settings can be overridden on a per user basis using the Active Directory Users and Computer snap-in (item 12.2.2). This can be done if only specific users have a legitimate need to send to large recipient lists. (It can also be used to create special accounts to which ordinary users can send mail requesting that the message be forwarded to a given, large recipient list. This effectively makes messages sent to large recipient lists moderated.)

1.2.6 Filter Recipients who are not in Directory

This controls whether email messages that are addressed to individuals that are not in the Domain's active directory be filtered immediately.

Clear this checkbox (that is this filter should not be applied). At first glance it would seem reasonable to want to immediately filter messages to recipients who are not in Active Directory since mail accounts are, in fact, stored in Active Directory. However, this feature can be used by external entities to determine whether a particular user exists in the Active Directory domain. By monitoring whether or not messages are filtered, an external entity could build a list of known accounts on the system. To prevent this disclosure of information, we recommend that this feature not be employed.

If this feature is enabled, administrators should ensure that user email addresses are not the same as the associated Windows account names (by default they are the same) because then learning someone's email address also gives someone the username of a system account. In fact, this is a good course of action to take in general, but especially important when external parties can easily learn the email addresses. If the administrator has made sure that email addresses are different from the usernames, then the threat from this feature is greatly mitigated and administrators may wish to consider enabling it in order to reduce the server load due to the generation of non-delivery reports. Also, the administrator should read KB 842851, which discusses enabling a tarpitting feature in Windows Server 2003. The tarpitting feature makes harvesting valid email addresses less attractive. Administrators should monitor account activity for suspicious signs, since learning the identity of an active account provides malicious outsiders with information they can use to begin mounting attacks.

Note that unless the “Apply recipient filter” checkbox (item 11.3.2) is selected on the SMTP Virtual Server, this control will not take effect on that virtual server regardless of the setting of this control.

1.2.7 Recipients

This field is used to filter messages based on their stated recipients. If this filter is enabled, any messages sent to a filtered recipient will be dropped early in the transmission process. (If multiple recipients are specified in a message, the message will still be delivered to recipients that are not listed in the filter regardless of whether other recipients match the filter.) Wildcards may be used to block more than one email address at a time.

List undesirable recipients in this field. The setting of this field will be entirely organization dependent. Some organizations may wish to create filters to prevent email from being sent to certain external addresses. (Note, however, that this will only work if internal users do not authenticate to the server - see below.) Other organizations may not need any recipient filters.

Note that unless the “Apply recipient filter” checkbox (item 11.3.2) is selected on the SMTP Virtual Server, this control will not take effect on that virtual server regardless of the setting of this control. Note also that these filters will only be applied to anonymous connections - users and computers that are able to authenticate will not be filtered using these rules.

1.2.8 Senders

This field is used to filter messages that have been sent from specified addresses. If this filter is enabled, email whose "From" field matches an entry in the list of prohibited senders will be dropped. Wildcards may be used to block more than one sender address at a time. This field can be used to block addresses of known spammers (although block list services, item 1.2.2, are generally used for this purpose), as well as blocking offensive or abusive users from sending email to the organization.

Enter the email addresses of malicious senders in this field. This field is useful when blocking individual users. (The global deny lists, item 1.2.1, can be used to block entire domains, but this would be overkill for a single abusive user.) All mail from this sender will be discarded when it arrives regardless of the intended destination.

Note that unless the "Apply sender filter" checkbox (item 11.3.2) is selected on the SMTP Virtual Server, this control will not take effect on that virtual server regardless of the setting of this control.

1.2.9 Archive Filtered Messages

This feature enables the archiving of all messages that are blocked by the sender filter. This file can be used to recover messages that might have been inappropriately filtered and to analyze the blocked messages that the organization has received. Note that this archive only contains messages that have been blocked by the sender filter.

Enable this feature. It provides a backup copy of filtered messages so that if, at a later point, it becomes apparent that a message was incorrectly filtered, that message can be recovered. Note that if this feature is enabled, the archive file can grow very rapidly, especially if the organization is targeted with a large amount of spam. For this reason, archives should be backed up and removed from the server on a regular basis. Depending on the amount of mail that is blocked by the sender filter, this may need to be done from every few days to more than once a day. Unless long term analysis of blocked messages is planned, archives will probably not need to be kept for more than a few weeks.

Failure to implement this feature may result in the permanent loss of messages that were incorrectly filtered. It will also deny administrators of data that can be used to determine patterns of spam and other abuse.

Note that unless the "Apply sender filter" checkbox (item 11.3.2) is selected on the SMTP Virtual Server, this control will not take effect on that virtual server regardless of the setting of this control.

1.2.10 Filter Messages with Blank Sender

This feature causes all messages received by the Exchange server that have a blank sender field to be blocked.

Enable this feature. Anonymous email (messages with blank sender fields) cannot be replied to and are almost invariably spam.

Failure to implement this feature will almost certainly result in an increase in the amount of spam individual users must deal with.

Note that unless the "Apply sender filter" checkbox (item 11.3.2) is selected on the SMTP Virtual Server, this control will not take effect on that virtual server regardless of the setting of this control.

1.2.11 Drop Connection if Address Matches Filter

This control allows the administrator to specify that any inbound connections from an address that is to be filtered should be immediately dropped. This is the most efficient way to handle the message from the Exchange server's perspective since it results in spending the least amount of time handling a message that will never be delivered anyway. This control only applies to messages that are blocked by the sender filter.

Enable this feature. In addition to minimizing the server resources taken up by a message that will be filtered, if the other party has other messages to send, it must re-initiate the SMTP connection to start sending the next message (as opposed to simply continuing the current connection). This will slow down the rate at which this blocked sender is able to send messages to the server, further reducing the resources that are likely to be taken up by the Exchange server.

If this setting is not selected, messages from unwanted senders will take up more processing power than necessary. If filtered connections are not dropped, select the “Accept messages without notifying sender of filtering” checkbox (item 1.2.12) to avoid giving spammers any clues as to the liveness of the destination to which they are sending their messages.

Note that unless the “Apply sender filter” checkbox (item 11.3.2) is selected on the SMTP Virtual Server, this control will not take effect on that virtual server regardless of the setting of this control.

1.2.12 Accept Messages Without Notifying Sender of Filtering

Selecting this feature prevents messages from being sent to blocked senders informing them that their message has been filtered. If it is not selected, senders that are filtered will receive a delivery failure notification that alerts them that the organization is filtering mail sent by them. This control only applies to messages that are blocked by the sender filter. This checkbox is not available if “Drop connection if address matches filter” (item 1.2.11) is selected.

If “Drop connection if address matches filter” (item 1.2.11) is not selected, enable this feature (Accept Messages Without Notifying Sender of Filtering). Often spammers will use delivery failure messages to determine the liveness of a particular target. While there is a chance that the blocked sender may stop sending mail on the grounds that it will not be delivered, it is more likely that the sender will simply note that it is sending to a live domain and continue to attempt to send messages, possibly using a different email address which might not be blocked. Moreover, the creation and sending of filtering notifications consumes server resources. By not creating these messages, one reduces the resources the server must use to deal with the blocked messages. That said, the most efficient course of action is simply to drop the connection from the banned sender immediately - (item 1.2.11).

If “Drop connection if address matches filter” (item 1.2.11) is not selected, failure to implement this recommendation will result in unnecessary use of server resources while alerting spammers and other malicious entities of server liveness. (Note that “Drop connection if address matches filter” (item 1.2.11) represents an even more efficient and secure configuration.)

Note that unless the “Apply sender filter” checkbox (item 11.3.2) is selected on the SMTP Virtual Server, this control will not take effect on that virtual server regardless of the setting of this control.

1.3 Enabled Services

1.3.1 Enable Outlook Mobile Access

These checkboxes control Outlook Mobile Access (OMA) on Exchange. While ActiveSync (item 1.3.2) is used simply to synchronize information between mobile devices and Exchange, OMA is used to provide an Outlook-like interface for mobile devices that offers many of the features of using Outlook itself. This is done by sending markup information (commonly HTML although other markup languages are supported) that creates a display that can then be utilized like a web page in any web compatible mobile device. The "Enable Outlook Mobile Access" checkbox enables this feature on Exchange. If OMA is enabled, the "Enable unsupported devices" checkbox will become available. Setting this control will cause Exchange to attempt to provide OMA pages (that is to say, pages formatted with markup tags) to any requesting devices and not just the devices it knows are capable of displaying OMA pages correctly.

Disable OMA. This is done to reduce the vectors through which connections may be made to your Exchange network. If this recommendation is followed the OMA Virtual Directory should be deleted from IIS (IIS Manager → [server] → Web Sites → Default Web Site → OMA). (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) This will prevent users from accessing the unused virtual directory and obviate the need to configure it (items 2.1.8, 2.4.7, 2.4.8). If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380.

If OMA is enabled on your system, the OMA connectivity will need to be monitored for suspicious activity. Moreover, a policy will need to be developed and users will need to be educated regarding secure practices with mobile devices. If OMA is enabled, allowing Exchange to connect to unsupported devices does not significantly increase the risk to the system and can be enabled or disabled as desired. If OMA is enabled be sure to use a secure channel for communication (item 2.1.12).

1.3.2 Enable User Initiated Synchronization

These checkboxes control the ActiveSync service on Exchange. ActiveSync is used to synchronize e-mail, calendaring, and contact information between the Exchange server and Windows supporting mobile devices, such as PDAs. To enable ActiveSync, "Enable user initiated synchronization" must be selected. If ActiveSync is enabled, the "Enable up-to-date notifications" checkbox becomes available and may be selected to send out alerts to the user's mobile device when new mail has arrived. (Normally, all downloads are strictly initiated by the mobile device itself.) If up-to-date notifications are enabled, the "Enable notifications to user specified SMTP addresses" checkbox becomes available and may be selected to allow individual users to select their own wireless service provider for up-to-date notifications.

Disable ActiveSync by clearing "Enable user initiated synchronization". This recommendation is made to restrict connectivity methods to the server. This can help reduce attack vectors as well as reduce the ways in which viruses may be introduced to your network. In addition, disabling ActiveSync provides (admittedly limited) control over how easily messages and other posts are allowed to leave your exchange network. If this recommendation is followed the Microsoft-Server-ActiveSync Virtual Directory should be deleted from IIS (IIS Manager → [server] → Web Sites → Default Web Site → Microsoft-Server-ActiveSync). (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) This will prevent users from accessing the unused virtual directory and obviate the need to configure it (items 2.1.7, 2.4.5, 2.4.6). If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If this recommendation is not followed, users with unauthorized physical access to a computer would be able to download data (for example, email and calendar) from that computer to a portable device.

If you choose to allow ActiveSync on your network, you should create policies for users regarding the correct and secure use of their handheld devices and develop ways to enforce them. Connectivity should be monitored to reduce the risk of malicious connections. The "Enable up-to-date notifications" feature does not significantly increase one's exposure if ActiveSync is enabled, so this feature may be set as desired. If ActiveSync is enabled, be sure to use a secure channel for communication (item 2.1.11).

1.4 Logging

1.4.1 Archive Filtered Messages See item 1.2.9

2. HTTP

2.1 Authentication

2.1.1 Certificate Wizard

This button starts a wizard to install a certificate to be used by this HTTP Virtual Server. Server certificates are required for many security features in Exchange, and without them the server cannot engage in many forms of secure communication. The wizard can guide the administrator through the process of requesting a new certificate or of importing an existing certificate. Certificates must be manually installed on each virtual server. This means that installing a certificate on one HTTP Virtual Server does not give other HTTP Virtual Servers (or the virtual servers of any other protocol) access to this certificate. However, once a certificate is installed on one virtual server, any other virtual server (regardless of protocol used) may easily be configured to use this certificate by selecting "Assign an existing certificate" in the first page of the Wizard.

Install certificates on this virtual server. Without them, many other recommendations in this document concerning secure communication will not be possible. For highest security assurance, each virtual server should have its own certificate that it does not share with other servers. This reduces the damage due to server compromises and provides per-server identification.

Failure to implement this recommendation makes it virtually impossible to secure Exchange's communications. Use of any virtual server that has not been given a certificate should be considered a highly insecure action.

2.1.2 Enable Anonymous Access

This controls whether users may connect to this virtual server without providing a username and password. If this is not selected, access will only be allowed to users who can be authenticated.

If any changes are made to this panel, the IIS manager will give the administrator the option of having these changes applied to all of the virtual directories residing on this virtual server. We do not recommend propagating the changes because doing so would overwrite the directory-by-directory recommendations that are provided elsewhere in this document. If the changes are not immediately propagated to the virtual directories, changing this setting will only affect virtual directories that are created after the change and, as such, should not affect Exchange security. Therefore, no recommendation is provided regarding this control's setting. (However, we do recommend that if its setting is changed, the changes should not be propagated to the individual virtual directories.)

2.1.3 Username/Password

This field specifies the user identity used by processes associated with anonymous HTTP sessions. Unlike other protocols, HTTP servers will often create separate processes that are associated with individual user requests, hence the need to specify a user identity under which such processes should run.

By default, the IUSR_<computer-name> user identity is specified (where <computer-name> is the name of the computer on which the virtual server is running). This account is specifically created for processes that are initiated by anonymous Internet users and has reduced access to the computer as a whole. The IUSR_<computer-name> account should be used.

Different user identities can be specified in this field, but they should not have greater access to the computer than the IUSR_<computer-name> account. Using an account with greater privileges puts the local computer at risk since it increases the capabilities of malicious anonymous users.

2.1.4 Authenticated Access

The Default Web site is the virtual server on which all Exchange virtual directories reside. This feature controls the authentication method used to connect to this virtual server and its virtual directories. Options include anonymous, Basic authentication (with clear text password), Digest authentication, and Integrated Windows Authentication.

Leave the default value of Integrated Windows Authentication only. Anonymous access provides for no access control of this virtual server, basic authentication transmits the password in the clear and risks exposure, and the other methods are not recommended by Microsoft for this control.

Failure to configure this as per the recommendation may result in unrestricted access to this virtual server, passwords being sent in the clear, and/or the inability to correctly authenticate, depending on which change is made.

Note that Integrated Windows Authentication cannot be used through front-end servers. As such, the recommended value effectively prohibits access to this virtual server through front-end servers.

Note also that if any changes are made to this panel, the IIS manager will give the administrator the option of having these changes applied to all the virtual directories residing on this virtual server. In general, this option should not be exercised as it will likely overwrite configuration information that has been made on the virtual directory level.

2.1.5 Authentication Method to Access Exchange Virtual Directory

The Exchange Virtual Directory is used to allow web access to user mail accounts. This is called Outlook Web Access (OWA) since it is designed to provide much of the same functionality provided by using an Outlook client, but through a web browser. This feature controls the authentication method used to connect to this virtual directory. Options include anonymous, Basic authentication (with clear text password), Digest authentication, and Integrated Windows Authentication.

OWA should not be used. As such, the recommendation is to delete the directory. The virtual directory should be deleted from IIS (IIS Manager → [server] → Web Sites → Default Web Site → Exchange). (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If OWA must be used, use the default authentication methods of Integrated Windows Authentication and Basic authentication. Since basic authentication is required for use by some client applications, it is vitally important that a secure channel be used with this control (item 2.1.14). Failure to do this will result in the password being sent in the clear.

If this directory is not deleted and if different authentication mechanisms are selected this may result in uncontrolled access to the directory or the inability of clients to correctly authenticate. Even if the directory is correctly configured there is always the possibility that a vulnerability may allow access through this directory. As such, the ideal course of action is to delete the directory and preclude any possibility of attack through this vector.

2.1.6 Authentication Method to Access Exadmin Virtual Directory

The Exadmin Virtual Directory is used by the Exchange System Manager to access mailboxes and public folders. As such, it is a required part of the Exchange application. This feature controls the authentication method used to connect to this virtual directory. Options include Anonymous, Basic authentication (with clear text password), Digest authentication, and Integrated Windows Authentication.

Leave this feature at the default value of Integrated Windows Authentication only. Anonymous access provides for no access control of this virtual directory, Basic authentication transmits the password in the clear, and the other methods are not recommended by Microsoft for this control.

Failure to configure this as per the recommendations may result in unrestricted access to this directory, passwords being sent in the clear, and/or the inability to correctly authenticate, depending on which change is made.

Note that Integrated Windows Authentication cannot be used across front-end servers. As such, the recommended value effectively prohibits access to this virtual directory across front-end servers.

Note also that this control cannot be set through the Exchange System Manager panel. It can be changed via the IIS manager, but this should not be done.

2.1.7 Authentication Method to Access ActiveSync Virtual Directory

The Microsoft-Server-ActiveSync Virtual Directory is used to provide access to ActiveSync services which are used to synchronize Exchange calendars and mailboxes with those on Exchange compliant portable devices such as PDAs. This feature controls the authentication method used to connect to this virtual directory. Options include Anonymous, Basic authentication (with clear text password), Digest authentication, and Integrated Windows Authentication.

If ActiveSync is not being used, the virtual directory should be deleted as per the recommendation in item 1.3.2. Doing this removes any need for further configuration of this control. If ActiveSync is used, leave this feature at the default value of Basic authentication. Anonymous access provides for no access control and ActiveSync cannot perform Integrated Windows Authentication. Note that it is vitally important that a secure channel be enabled for this virtual directory. (See item 2.1.11). Failure to do this will result in passwords being transmitted in the clear.

If this virtual directory is not deleted (against the recommendation in item 1.3.2) and anonymous access is permitted, then users will be able to connect to the server without any form of authentication. Enabling Integrated Windows Authentication or Digest authentication is unlikely to have much of an effect since ActiveSync cannot use these protocols.

2.1.8 Authentication Method to Access Outlook Mobile Access Virtual Directory

The OMA Virtual Directory is used to provide access to Outlook Mobile Access services which are used to provide an interface between Exchange and Exchange compliant portable devices such as certain types of cell phones and PDAs. This feature controls the authentication method used to connect to this virtual directory. Options include Anonymous, Basic authentication (with clear text password), Digest authentication, and Integrated Windows Authentication.

If OMA is not being used, the virtual directory should be deleted as per the recommendation in item 1.3.1. If OMA is used, leave this feature at the default value of Basic authentication. Anonymous access provides for no access control and OMA cannot perform Integrated Windows Authentication. Note that it is vitally important that a secure channel be enabled for this virtual directory. (See item 2.1.12) Failure to do this will result in passwords being transmitted in the clear.

If this virtual directory is not deleted (against the recommendation in item 1.3.1) and Anonymous access is permitted, then users will be able to connect to the server without any form of authentication. Enabling Integrated Windows Authentication or Digest authentication is unlikely to have much of an effect since OMA cannot use these protocols.

2.1.9 Enable Forms Based Authentication

This controls how authentication to Outlook Web Access (OWA) is handled. When the “Enable forms based authentication” checkbox is selected, authenticating users are sent a form through which they log in to Exchange. This form stores username and password information in browser cookies (so cookies must be enabled in the client browser for this to work). These cookies persist throughout the OWA session after which they are destroyed. (Cookies may also timeout automatically after a period of inactivity.) Once the cookie has been destroyed, the user must re-authenticate to Exchange before they will be permitted any further interaction with the server. If forms based authentication is not used, standard HTTP authentication will be used instead. In HTTP based authentication, the username and password will be retained as long as the browser remains running. Even after the user logs out from Exchange, if anyone using the same browser session attempts to reconnect to the Exchange server, the cached username and password will automatically authenticate them to the server. The username and password will remained cached until the browser application is terminated.

Enable forms based authentication. Doing so makes it less likely that an unauthorized user will be able to hijack someone's account using the web browser that a legitimate user last utilized. This is because all authentication credentials are destroyed as soon as the session ends. Moreover, some environments, such as Internet kiosks, may not allow users to close browsers giving users no way of protecting their account. If forms based authentication is not used, credentials remain for a much longer period of time, giving an unauthorized user a greater window of opportunity.

Failing to use forms based authentication places a higher burden on client users in terms of protecting access to their account. Specifically, in addition to logging out, they will also need to make sure the browser program is terminated before they leave the terminal where they accessed their account. If they always remember to do so, they face no additional risk, however it is one more step that is required of the user. Enabling forms based authentication means the user only needs to log out from their Exchange account to protect themselves against hijacking.

Regardless of the setting of this control, if OWA is to be enabled, users will need to be educated as to good security procedures (including the appropriate signoff behaviors as described above) when browsing Exchange mail using web accounts.

2.1.10 Authentication Method to Access Public Virtual Directory

The Public Virtual Directory is used to provide access to public folders. This feature controls the authentication method used to connect to this virtual directory. Options include Anonymous, Basic authentication (with clear text password), Digest authentication, and Integrated Windows Authentication.

If public folders will not be used on this Exchange server the virtual directory should be deleted to remove it as a potential attack vector. This should be done using the IIS manager (IIS Manager → [server] → Web Sites → Default Web Site → Public). (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If public folders must be used, leave this feature at the default value of Integrated Windows Authentication and Basic authentication.

If this directory is not deleted and if different authentication mechanisms are selected this may result in uncontrolled access to the directory or the inability of clients to correctly authenticate. Even if the directory is correctly configured there is always the possibility that a vulnerability may allow access through this directory. As such, the ideal course of action is to delete the directory and preclude any possibility of attack through this vector.

2.1.11 Require Secure Channel/Require 128 bit encryption to Microsoft-Server-ActiveSync Virtual Directory

This controls whether client machines should be forced to use secure channels to communicate with this virtual directory. If this feature is enabled, clients will only be able to communicate with the directory if they are capable of supporting secure communication with the server. If secure channels are required, the server can also require the channel to be strongly secure by requiring 128 bit encryption.

Ideally, the Microsoft-ActiveSync Virtual Directory will be deleted. However, if ActiveSync is to be used, require secure channels and 128-bit encryption. The use of secure communication prevents eavesdroppers from reading or modifying communications between servers and clients.

If a particular client application in your organization does not support secure communication, this feature will need to be disabled or that client will be unable to connect to the server. However, it is strongly advised that any such clients be upgraded. Failure to enable this feature gives attackers more opportunities to read and modify communication between the client and server.

Requiring secure channels but not requiring 128-bit encryption will still protect most messages from eavesdropping and modification, but makes it slightly more likely that a determined and well equipped attacker may be able to overcome the channel security. This setting is particularly important for ActiveSync since ActiveSync uses basic authentication which sends password information over the network in the clear. Failure to enable a secure channel will likely result in password disclosure. Note, however, that the ideal solution is simply to delete this virtual directory which eliminates the possibility of attack through this vector.

2.1.12 Require Secure Channel (SSL) and Require 128 bit encryption to OMA Virtual Directory

This controls whether client machines should be forced to use secure channels to communicate with this virtual directory. If this feature is enabled, clients will only be able to communicate with the directory if they are capable of supporting secure communication with the server. If secure channels are required, the server can also require the channel to be strongly secure by requiring 128 bit encryption.

Ideally, the OMA Virtual Directory will be deleted. However, if OMA is to be used, require secure channels and 128-bit encryption. The use of secure communication prevents eavesdroppers from reading or modifying communications between servers and clients.

If a particular client application in your organization does not support secure communication, this feature will need to be disabled or that client will be unable to connect to the server. However, it is strongly advised that any such clients be upgraded. Failure to enable this feature gives attackers more opportunities to read and modify communication between the client and server. Requiring secure channels but not requiring 128-bit encryption will still protect most messages from eavesdropping and modification, but makes it slightly more likely that a determined and well equipped attacker may be able to overcome the channel security. This setting is particularly important for OMA since OMA uses basic authentication which sends password information over the network in the clear. Failure to enable a secure channel will likely result in password disclosure. Note, however, that the ideal solution is simply to delete this virtual directory which eliminates the possibility of attack through this vector.

2.1.13 Require Secure Channel (SSL) and Require 128 bit encryption to Public Virtual Directory

This controls whether client machines should be forced to use secure channels to communicate with this virtual directory. If this feature is enabled, clients will only be able to communicate with the directory if they are capable of supporting secure communication with the server. If secure channels are required, the server can also require the channel to be strongly secure by requiring 128 bit encryption.

Ideally, the Public Virtual Directory should be deleted. However, if public folders are to be used, require secure channels and 128-bit encryption. The use of secure communication prevents eavesdroppers from reading or modifying communications between servers and clients. Note that, if public folders are used, and if a front-end/back-end topology is employed by the system, the back-end server should NOT require a secure channel for this virtual directory. This is because the communication between front-end and back-end servers does not support SSL. As a result, if the back-end server's Public Virtual Directory is configured to require a secure channel, this will prevent the use of this virtual directory via the front-end server. However, all front-end servers and standalone Exchange servers should require a secure channel.

If a particular client application in your organization does not support secure communication, this feature will need to be disabled or that client will be unable to connect to the server. However, it is strongly advised that any such clients be upgraded. Failure to enable this feature gives attackers more opportunities to read and modify communication between the client and server. Requiring secure channels but not requiring 128-bit encryption will still protect most messages from eavesdropping and modification, but makes it slightly more likely that a determined and well equipped attacker may be able to overcome the channel security. This setting is particularly important for the Public Virtual Directory since it uses basic authentication which sends password information over the network in the clear. Failure to enable a secure channel will likely result in password disclosure. Note, however, that the ideal solution is to delete this virtual directory which eliminates the possibility of attack through this vector.

2.1.14 Require Secure Channel (SSL) and Require 128 bit encryption to Exchange Virtual Directory

This controls whether client machines should be forced to use secure channels to communicate with this virtual directory. If this feature is enabled, clients will only be able to communicate with the directory if they are capable of supporting secure communication with the server. If secure channels are required, the server can also require the channel to be strongly secure by requiring 128 bit encryption.

Ideally, the Exchange Virtual Directory should be deleted. However, if OWA is to be used, require secure channels and 128-bit encryption. The use of secure communication prevents eavesdroppers from reading or modifying communications between servers and clients. Note that, if OWA is used, and if a front-end/back-end topology is employed by the system, the back-end server should NOT require a secure channel for this virtual directory. This is because the communication between front-end and back-end servers does not support SSL. As a result, if the back-end server's Exchange Virtual Directory is configured to require a secure channel, this will prevent the use of this virtual directory via the front-end server. However, all front-end servers and standalone Exchange servers should require a secure channel.

If a particular client application in your organization does not support secure communication, this feature will need to be disabled or that client will be unable to connect to the server. However, it is strongly advised that any such clients be upgraded. Failure to enable this feature gives attackers more opportunities to read and modify communication between the client and server. Requiring secure channels but not requiring 128-bit encryption will still protect most messages from eavesdropping and modification, but makes it slightly more likely that a determined and well equipped attacker may be able to overcome the channel security. This setting is particularly important for the Exchange Virtual Directory since it uses basic authentication which sends password information over the network in the clear. Failure to enable a secure channel will likely result in password disclosure. Note, however, that the ideal solution is to delete this virtual directory which eliminates the possibility of attack through this vector.

2.1.15 Require Secure Channel (SSL) and Require 128 bit encryption to Exadmin Virtual Directory

The Exadmin Virtual Directory is used by the Exchange System Manager to access mailboxes and public folders. As such, it is a required part of the Exchange application. This feature controls the authentication method used to connect to this virtual directory. Options include Anonymous, Basic authentication (with clear text password), Digest authentication, and Integrated Windows Authentication.

This item controls whether client machines should be forced to use secure channels to communicate with this virtual directory. If this feature is enabled, clients will only be able to communicate with the directory if they are capable of supporting secure communication with the server. If secure channels are required, the server can also require the channel to be strongly secure by requiring 128 bit encryption.

The services that use the Exadmin Virtual Directory do not support the use of secure channels. As such, secure channels should not be required on this virtual directory. Failure to follow this recommendation will effectively disable Exchanges mail and public folder functionality.

2.2 Connections

2.2.1 TCP Port/SSL Port

This controls the ports to which the HTTP and HTTPS servers (regular and secure HTTP) bind. The standard port for regular HTTP connections is 80 and the standard port for HTTPS connections is 443. If different ports are used, clients will need to be explicitly configured to use the non-standard ports.

Use the standard ports (that is, the default). Traffic on these ports should be carefully monitored for signs of suspicious activity.

Changing the ports to non-standard values can provide some limited protection against automated attacks since these attacks will not connect to the custom port, and therefore be nullified. However, making this modification introduces a large amount of complexity for the system administrator: If clients are to connect to the server, they will all need to be configured to use the selected non-standard ports for their communication. In addition, if at some later point, additional network services are added to the Exchange server, care must be taken that the ports used by these services do not conflict with the non-standard ports chosen here. Additionally, since this feature is set for the entire IIS Default Web Site and not just Exchange, factors other than the needs of Exchange may need to be considered. For example, if this IIS server is also running an open, public web site, changing the ports would prevent the general public from reaching the information on the site. Finally, if clients outside the firewall or proxy server are to be allowed to connect to the server, or if automated network monitoring tools are employed, these services will all need to be configured to use the non-standard ports for HTTP/HTTPS traffic. Even if ports are changed, however, a determined attacker may still be able to determine which ports are used for the HTTP and HTTPS protocols by performing a comprehensive port scan, although doing so should be detectable by most network monitoring tools. Since changing the ports introduces a large amount of complexity for a relatively small gain, the standard ports should be used.

2.3 Delivery Restrictions

2.3.1 Address/Domain Name Restrictions

This controls which IP addresses are allowed to connect to this virtual server. The control can be set to either allow all computers to connect except for a specified few, or to deny all computers except for a specified few. In addition to individual IP addresses, subnet masks may be used to specify groups of addresses. Note that if your network uses DHCP to dynamically assign IP addresses to client computers it will not be possible to reliably specify one computer out of the group serviced by a given DHCP server since that computer's IP address may change regularly. (If all computers serviced by a DHCP server should have the same treatment then there is no problem since one or more subnet masks can then be used to cover the DHCP server's entire address pool, ensuring that regardless of the address a computer is assigned, it will have the same treatment in this panel.)

The default setting is “Granted to all with no exceptions”. This is likely the correct configuration for most enterprises. However, there may be cases in which one would wish to exclude a list of addresses or even only allow access to a specific list of hosts. As such, no specific recommendation is given here.

2.4 Directory Access

2.4.1 Execute Permissions for OWA Virtual Directory

The Exchange Virtual Directory is used to allow web access to user mail accounts. This is called Outlook Web Access (OWA) since it is designed to provide much of the same functionality provided by using an Outlook client, but through a web browser. This control allows the administrator to specify whether scripts and/or executables may be run on this virtual directory.
Delete the Exchange Virtual Directory if the service is not to be used. If this recommendation is followed the Exchange Virtual Directory should be deleted from IIS (IIS Manager → [server] → Web Sites → Default Web Site → Exchange). (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If OWA is to be used set the value of this control to "None", which is the default value.

Allowing scripts or executables to be run on this virtual directory opens up a powerful and unnecessary capability. Scripts on servers are a frequent cause of server compromises. Moreover, since this virtual directory is the primary interface between Exchange and the web, it is particularly at risk of compromise. As a result, every attempt to minimize attack vectors especially via scripts and executables running on the server, should be minimized.

2.4.2 Access Control to OWA Virtual Directory

The Exchange Virtual Directory is used to allow web access to user mail accounts. This is called Outlook Web Access (OWA) since it is designed to provide much of the same functionality provided by using an Outlook client, but through a web browser. This control determines whether users will have read, write, script source access, and/or directory browsing capabilities to this virtual directory.

Delete the Exchange Virtual Directory if the service is not to be used. If this recommendation is followed the Exchange Virtual Directory should be deleted from IIS (IIS Manager → [server] → Web Sites → Default Web Site → Exchange) (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If OWA is to be used, all four of read, write, script source access, and directory browsing will need to be selected since these are required for the proper functioning of OWA. This is the default setting for OWA Access Control.

2.4.3 Access Controls to Exadmin Virtual Directory

The Exadmin Virtual Directory is used by the Exchange System Manager to access mailboxes and public folders. As such, it is a required part of the Exchange application. This control allows the administrator to specify the combination of read, write, script source access, and directory browsing capabilities allowed to this virtual directory.

Allow all four of these capabilities on this virtual directory (select all four checkboxes). The Exchange System Manager is a central part of the Exchange application and without these capabilities it will be unable to function properly. Granting all four capabilities is the default value of this control. In fact, Exchange does not allow this control to be altered through the Exchange management console, but it is possible to change the value of this setting using IIS (IIS Manager → [server] → Web Sites → Default Web Site → Exadmin → Properties → Virtual Directory Tab). However, this should not be done.

Changing this setting, while reducing the capabilities of applications accessing this directory, will inhibit the proper functioning of the Exchange Systems Manager. Specifically, the Exchange System Manager uses this virtual directory to administer public folders via the web and without it this activity will not be possible.

2.4.4 Execute Permissions to access Exadmin Virtual Directory

The Exadmin Virtual Directory is used by the Exchange System Manager to access mailboxes and public folders. As such, it is a required part of the Exchange application. This control allows the administrator to specify whether scripts and/or executables may be run on this virtual directory.

Scripts and executables should be denied the ability to run on this directory (set to None). This is the default value of this control. In fact, Exchange does not allow this control to be altered through the Exchange management console, but it is possible to change the value of this setting using IIS (IIS Manager → [server] → Web Sites → Default Web Site → Exadmin → Properties → Virtual Directory Tab). However, it is highly recommended that this not be done.

Changing this setting grants powerful and unnecessary capabilities to users. Since the Exchange System Manager is the only entity that should be interfacing with this virtual directory, and since the virtual directory provides all of the capabilities the manager needs by default, there is no reason to grant any additional privileges.

2.4.5 Access Control to ActiveSync Virtual Directory

The Microsoft-Server-ActiveSync Virtual Directory is used to provide access to ActiveSync services. These services are used to synchronize Exchange calendars and mailboxes with those on Exchange compliant portable devices such as PDAs. This control governs the read, write, script source access, and directory browsing capabilities that users and computers will have to this directory. For ActiveSync to be enabled, read must be enabled.

Delete the Microsoft-Server-ActiveSync Virtual Directory if the service is not to be used. If this recommendation is followed the Microsoft-Server-ActiveSync Virtual Directory should be deleted from IIS (IIS Manager → [server] → Web Sites → Default Web Site → Microsoft-Server-ActiveSync). (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If ActiveSync is to be used, only read permission should be granted to the virtual directory.

If ActiveSync is used and the directory is not deleted, selecting more than read access grants unnecessary levels of access to this directory. Since ActiveSync only needs read access in order to function, there is no reason to grant further permissions. Ideally, however, the virtual directory should be deleted completely eliminating it as a possible attack vector.

2.4.6 Execute Permissions to access ActiveSync Virtual Directory

The Microsoft-Server-ActiveSync Virtual Directory is used to provide access to ActiveSync services, which are used to synchronize Exchange calendars and mailboxes with those on Exchange compliant portable devices such as PDAs. This control allows an administrator to specify whether scripts and/or executables may be run on this virtual server. For ActiveSync to function, “Scripts and Executables” must be selected.

If you are not using ActiveSync, delete the virtual directory. If this recommendation is followed the Microsoft-Server-ActiveSync Virtual Directory should be deleted from IIS (IIS Manager → [server] → Web Sites → Default Web Site → Microsoft-Server-ActiveSync). (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If you do plan on using ActiveSync, use the default value of allowing both scripts and executables to run as ActiveSync cannot function without these permissions.

Ideally, this virtual directory should be deleted to eliminate it as an attack vector. If ActiveSync is used this control cannot help lock it down since ActiveSync requires the highest level of access to operate.

2.4.7 Access Control to OMA Virtual Directory

The OMA Virtual Directory is used to provide access to Outlook Mobile Access services, which are used to provide an interface between Exchange and Exchange compliant portable devices such as certain types of cell phones and PDAs. This control governs read, write, script source access, and directory browsing capabilities that users and computers have to this directory. For OMA to be enabled, read must be enabled.

Ideally, the OMA Virtual Directory should be deleted from IIS (IIS Manager → [server] → Web Sites → Default Web Site → OMA). (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If OMA is to be used, only grant read access to this virtual directory.

If OMA is to be used and the virtual directory is not deleted, allowing more than read access grants users unnecessary access to this virtual directory OMA only needs read access to this directory so further access is not needed.

2.4.8 Execute Permissions to access OMA Virtual Directory

The OMA Virtual Directory is used to provide access to Outlook Mobile Access services, which are used to provide an interface between Exchange and Exchange compliant portable devices such as certain types of cell phones and PDAs. This control allows an administrator to specify whether scripts and/or executables may be run on this virtual directory. For OMA to function, "Scripts" must be selected.

If you are not using this, delete the OMA Virtual Directory from IIS (IIS Manager → [server] → Web Sites → Default Web Site → OMA). Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future. If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If OMA is to be used, the default value of allowing only scripts to run should be enabled.

Also, if using OMA and the virtual directory is not deleted, only allow scripts to run in this directory. OMA requires these scripts to operate, but increasing access by adding the ability to run executables is extremely dangerous since it gives powerful and unnecessary capabilities to users. Ideally, however, this virtual directory should be deleted to eliminate it as an attack vector.

2.4.9 Execute Permissions to access Public Virtual Directory

The Public Virtual Directory is used to provide access to public folders. This control allows administrators to specify whether scripts and/or executables may be run on this virtual directory.

Delete the Public Virtual Directory if public folders are to be used. If this recommendation is followed the Public Virtual Directory should be deleted from IIS (IIS Manager → [server] → Web Sites → Default Web Site → Public). (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If public folders are used, set the value of this control to none to prevent scripts or executables from running in this directory.

Note that this control may also be set using the IIS Manager (IIS Manager → [server] → Web Sites → Default Web Site → Public → Properties → Virtual Directory Tab). However, the settings specified in the Exchange console will override those set through the IIS Manager.

Allowing scripts or executables to be run on this virtual server opens up a powerful and unnecessary capability. Scripts on servers are a frequent cause of server compromises. As a result, every attempt to reduce attack vectors (especially via scripts and executables running on the server), should be made. Ideally, however, the directory should be deleted to preclude its use as an attack vector.

2.4.10 Access Controls to Public Virtual Directory

The Public Virtual Directory is used to provide access to public folders. This control determines whether users have read, write, script source access, and/or directory browsing capabilities on this virtual directory.

Delete the Public Virtual Directory if access to public folders is not needed. If this recommendation is followed the Public Virtual Directory should be deleted from IIS (IIS Manager → [server] → Web Sites → Default Web Site → Public). (Do not delete the directory from the Exchange System Manager as doing so will make it more difficult to restore the directory in the future.) (Note: deleting the Public Virtual Directory disables OWA access to public folders.) If necessary, the directory can be recreated by following the procedures outlined in Microsoft Knowledge Base article 883380. If public folders are to be used, read, write, script source access, and directory browsing rights must be granted for full functionality (these are the default settings).

Note that, if access to this virtual directory is denied, the (Exchange System Manager → Administrative Groups → [administrative group] → Folders → Public Folder → [folder] → Content Tab in the Exchange System Manager will not be able to display the contents of this public folder since it uses this virtual server to access the folder contents. The Exchange System Manager will still be able to administer public folders since this functionality is supported by the Exadmin Virtual Server.

Ideally, this virtual directory should be deleted to eliminate it as an attack vector. If public folders are used this control cannot help lock it down since public folders require full access to operate.

2.5 Logging

2.5.1 Enable Logging

This controls the creation and format of log files used to monitor activity on this virtual server. By default, these files will be stored in \WINDOWS\SYSTEM32\LOGFILES\W3SVCx (where x is a number used to distinguish between virtual servers in this organization). The drop-down menu is used to select the format of the log file. The properties button next to this dropdown displays configuration information specific to the type of log format selected, but usually has some control to indicate the log rotation schedule (that is, how often the old log file should be closed and a new log file should be started). The log files contain, among other things, the times that users connect, what articles they read, and from where they are connecting.

Enable logging. Doing so records attempts made to connect to this virtual server. In the case of an attack on the HTTP server, these logs could contain useful details regarding the time and nature of the attack. Due to the size of log files, the files should be regularly copied to external storage and deleted from the server to conserve memory. Log files should be retained for at least one month. The format of the log files is largely a matter of preference for the administrator.

Failure to enable this feature means that the administrator will not have access to information regarding HTTP connections should the HTTP servers come under attack.

It is required that all log files be written to separate partitions from the partitions used by the Exchange Stores. Exchange will dismount its stores if it detects that it has run out of disk space, resulting in a complete loss of Exchange services. To minimize the chance of this happening, log files should write to some other partition so that if the logs fill this partition it will not result in the failure of Exchange.

3. IMAP

3.1 Authentication

3.1.1 Authentication Method used to Access IMAP Virtual Directory

This controls the form of authentication used by clients attempting to connect to this virtual server. Possible settings are “Basic authentication” and “Simple Authentication and Security Layer”. If Basic authentication is selected, an additional checkbox allows the server to require that SSL/TLS encryption be used to protect the password (which would otherwise be sent in clear text). Simple Authentication and Security Layer is the Microsoft NTLM protocol.

Select Basic authentication and Require SSL/TLS in this panel. The use of SSL/TLS not only protects the username and password during authentication, but encrypts the mail messages as they are being transmitted, preventing eavesdroppers from reading messages.

Failure to implement SSL/TLS with Basic authentication is extremely insecure as it results in the user's identity and password being sent over the network in clear text. If this happens, network sniffers can easily collect these credentials. The use of NTLM (Simple Authentication and Security Layer checkbox), while it can protect the username and password during authentication, it does not provide encryption of message bodies, potentially allowing message content to be sniffed over the wire. Moreover, NTLM negotiates the details of how it secures the authentication and some possible selections are insecure. While it is possible to configure NTLM server to only use the more secure options, Exchange does not provide this capability. (To only permit the most secure form of the protocol, NTLM v2, set the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\LMCompatibility registry key to 5. See Microsoft Knowledge Base Article 239869.)

3.1.2 Certificate Wizard

This button starts a wizard to install a certificate to be used by this IMAP Virtual Server. Server certificates are required for many security features in Exchange, and without them the server cannot engage in many forms of secure communication. The wizard can guide the administrator through the process of requesting a new certificate or of importing an existing certificate. Certificates must be manually installed on each virtual server. This means that installing a certificate on one IMAP Virtual Server does not give other IMAP Virtual Servers (or virtual servers of any other protocols) access to this certificate. However, once a certificate is installed on one virtual server, any other virtual server (regardless of protocol used) may easily be configured to use this certificate by selecting “Assign an existing certificate” in the first page of the Wizard.

Install certificates on this virtual server. Without it, many other recommendations in this document concerning secure communication will be impossible. For highest security assurance, each virtual server should have its own certificate that it does not share with other servers. This reduces the damage due to server compromises and provides per-server identification.

Failure to implement this recommendation makes it virtually impossible to secure Exchange's communications. Use of any Exchange Virtual Server that has not been given a certificate should be considered a highly insecure action.

3.2 Connections
3.2.1 Require Secure Channel (SSL) and Require 128 bit encryption to IMAP Virtual Directory This controls whether client machines should be forced to use secure channels to communicate with this virtual server. If this feature is enabled, clients will only be able to send and download mail from the server if they are capable of supporting secure communication with the server. If secure channels are required, the server can also require the channel to be strongly secured by requiring 128 bit encryption. Require secure channels and 128-bit encryption. The use of secure communication prevents eavesdroppers from reading or modifying communications between servers and clients. While sensitive message bodies should be encrypted by the sender at the client, securing the communication between the server and the recipient provides an added level of security while also protecting sender and recipient information that cannot be encrypted by the sender. If a particular client application in your organization does not support secure communication, this feature will need to be disabled or that client will be unable to connect to the server. However, it is strongly advised that any such clients be upgraded. Failure to enable this feature gives attackers more opportunities to read and modify communication between the client and server. Requiring secure channels but not requiring 128-bit encryption will protect most messages from eavesdropping and modification, but makes it slightly more likely that a determined and well equipped attacker may be able to overcome the channel security.
3.2.2 Use SSL Connections This controls whether responses to meeting requests will use an encrypted channel when responding to servers. Meeting request emails contain links to an Outlook Web Access page that allow users to access most of the features of standard Outlook meeting requests over the web. The “Use SSL connections” field determines whether the download and subsequent server interaction via the downloaded form uses a secure channel. Select “Use SSL connections”. This protects user interactions with their calendar from being observed or modified by an eavesdropper. Failure to select this feature means that the use of OWA to handle meeting requests might be observed or modified by a third party. Note that for this to work, the HTTP “Exchange” Virtual Directory must exist and be configured to use secure channels (item 2.1.14).

3.2.3 TCP Port/SSL Port

This controls the ports that the IMAP Virtual Server binds to for regular and secure communications, respectively. The standard port for regular IMAP connections is 143, while the standard port for secure IMAP is 993. If ports other than these are to be used, clients will need to be explicitly configured to use the non-standard port.

Use standard ports for both protocols (that is, the default). Traffic on these ports should be carefully monitored for signs of suspicious activity.

Changing the port to a non-standard value can provide some limited protection against automated attacks since these attacks will not connect to the custom port, and therefore be nullified. However, making this modification introduces a large amount of complexity for the system administrator: If clients are to connect to the server, they will all need to be configured to use the selected non-standard ports for their communication. In addition, if at some later point, additional network services are added to the Exchange server, care must be taken that the ports used by these services do not conflict with the non-standard ports chosen here. Finally, if clients outside the firewall or proxy server are to be allowed to connect to the server, or if automated network monitoring tools are employed, these services will all need to be configured to use the non-standard ports for IMAP traffic. Even if ports are changed, however, a determined attacker may still be able to determine which port is used for the IMAP protocol by performing a comprehensive port scan, although doing so should be detectable by most network monitoring tools. Since changing the port introduces a large amount of complexity for a relatively small gain, the standard ports should be used.

3.2.4 Limit Number of Connections

This controls the maximum number of simultaneous connections allowed to the IMAP server. This can be used to throttle the use of the IMAP service if it begins to monopolize system resources. If the checkbox is cleared, no limits are applied to the number of connections. If the checkbox is selected, the text field on the right is used to specify the maximum number of connections.

The setting of this control is highly dependent on environment. In some cases, applying a limit to the number of IMAP connections is a useful way to control the network resources that are consumed by this protocol. This is especially likely if IMAP is not the primary role of this server. However, in other cases, applying a limit may be detrimental to the enterprise's mission because it may end up prohibiting legitimate IMAP use. This is particularly likely if the server is primarily intended to service IMAP requests and network resources are unlikely to be a limiting factor. Because of the range of scenarios that would affect the setting of this control, no recommendation is given here.

If no limit is provided, the administrator should closely monitor IMAP use to ensure that IMAP activity does not come at the expense of other protocols. If the administrator notices this happening, limiting the number of IMAP connections may help remediate the problem. (This symptom might, however simply indicate that the server's network resources are insufficient for its use.) Likewise, if a limit is applied, administrators should be careful that this limit is not so low that it needlessly limits legitimate IMAP connections.

3.2.5 Connection Time-out (Minutes)

This controls the number of minutes that an idle connection to the IMAP server will be maintained before being dropped by the server. It can be used to limit the number of idle connections that the server maintains.

The default value for this control is 30 minutes and cannot be reduced. It is recommended that this value not be increased as there is seldom a need to maintain an idle connection for longer.

If the timeout is increased then the server will end up needing to maintain more simultaneous connections than might be necessary. This is particularly significant if the number of simultaneous connections is limited as it means that unused connections count against this limit for a longer period of time. See item 3.2.4 for more information on limiting the number of simultaneous connections. It should be noted, however, that the server expends relatively few resources maintaining idle connections, so this control is unlikely to affect actual system resource usage except in very unusual circumstances.

3.3 Delivery Restrictions

3.3.1 Exclude or Limit Connections

This controls which IP addresses are allowed to connect to this virtual server to send or download messages. The control can be set to either allow all computers to connect except for a specified few, or to deny all computers except for a specified few. In addition to individual IP addresses, subnet masks may be used to specify groups of addresses. Note that if your network uses DHCP to dynamically assign IP addresses to client computers it will not be possible to reliably specify one computer out of the group serviced by a given DHCP server since that computer's IP address may change regularly. (If all computers serviced by a DHCP server should have the same treatment then there is no problem since one or more subnet masks can then be used to cover the DHCP server's entire address pool, ensuring that regardless of the address a computer is assigned, it will have the same treatment in this panel.)

Select “Only the list below” so that the administrator must explicitly specify which clients can connect to the IMAP Virtual Server. This significantly reduces the chance of unauthorized connections to the server and helps to further restrict network connectivity.
If “All except the list below” must be selected, administrators should monitor connectivity to the IMAP server to ensure that no suspicious connections are being made.

Note that, for proper IMAP operation, clusters must know about each other and their IPs need to be added. In addition, all back-end servers must be able to connect to all other back-end servers within their cluster as well as their corresponding front-end server. (That is, the IP address of these servers must be allowed in this panel). Likewise, all front-end servers must be able to connect to all legal clients as well as all back-end servers. (It may not be possible to know the IP addresses of all possible clients, but if your system uses a VPN, the IP assigned by the VPN can be used here instead. This scenario is recommended if external access is required.)

4. Mailbox Store

4.1 Authentication

4.1.1 Clients Support S/MIME Signatures

This checkbox is set if the email clients serviced by this store support the use of S/MIME signatures. If selected, Exchange will use S/MIME signatures to verify the source of messages.

All clients in the enterprise should be updated to support S/MIME. If there is a valid reason that S/MIME cannot be used, other client-specific mechanisms for message signing should be used. However, S/MIME is one of the most commonly supported mechanisms so its use is most likely to be compatible with recipient mail clients.

4.2 Backup/Restore

4.2.1 Time of Last Full Backup

This field lists the time of the most recent known full backup of this Mailbox Store. To be "known", the backup must either use the Windows Server Backup Wizard (Start → All Programs → Accessories → System Tools → Backup) or some other Windows/Exchange aware backup utility. Only Windows/Exchange aware backup utilities are recommended since other backup utilities may not be able to correctly restore the mailbox store. This field cannot be set - it is simply present to provide information. Select this field soon after each scheduled full backup to ensure that the backup completed successfully.

Full backups of the mailbox store should occur at least on a weekly basis. More frequent full backups will simplify any data restoration that may be necessary. Mailbox store backups should take place with or in addition to backups of the full server.

4.2.2 Time of Last Incremental Backup

This field lists the time of the most recent known incremental backup of this Mailbox Store. To be "known", the backup must either use the Windows Server Backup Wizard (Start → All Programs → Accessories → System Tools → Backup) or some other Windows/Exchange aware backup utility. Only Windows/Exchange aware backup utilities are recommended since other backup utilities may not be able to correctly restore the mailbox store. This field cannot be set - it is simply present to provide information. Select this field soon after each scheduled incremental backup to ensure that the backup completed successfully.

Incremental backups of the mailbox store should occur at least on a daily basis. More frequent backups can help to minimize the amount of work that is lost in case of a disk failure, but may result in a significant load on the server. The frequency of incremental backups should balance the importance of the information on the store against the ability of the server to handle incremental backups simultaneously with its other duties.

Mailbox store backups should take place with or in addition to backups of the full server.

4.2.3 Do Not Permanently Delete Mailboxes Until Backed Up

This control allows the administrator to prevent the permanent deletion of messages and mailboxes until after said entities have been backed up (using an Exchange aware backup utility such as the Microsoft Backup wizard from Start → All Programs → Accessories → System Tools → Backup). Exchange handles message deletion by keeping a count of the number of references to each individual message on the mailbox store. When a user "deletes" a message, they are actually just removing their reference to that message. Once all the references to a message have been deleted, that message is marked as being ready for final deletion. A maintenance process (scheduled by item 4.2.4) runs periodically and actually deletes all messages that are marked for final deletion. It is only at this point that the message is actually deleted from the store and its memory made available. This checkbox ensures that at least one backup has been run on the mailbox store before the message actually disappears.

Enable this feature since it significantly reduces the chance of accidental deletion of important information. It also means that all messages written to recipients who have accounts on this store will reside in backups even after deletion.

The only potential drawback to this is if the backup facility is inoperable for an extended period of time the size of the mailbox store may grow due to the fact that no messages are actually being deleted. Even under these circumstances, however, it is probably more important to wait until the backup service is restored then risk permanently losing important information.

4.2.4 Maintenance Interval

This controls when and for how long the Exchange Mailbox Store maintenance service will run. The maintenance service performs a number of activities including final deletion of files (as governed by other policy controls) and checking mailbox storage limits against policy. The administrator is given a very fine degree of control of when the maintenance service runs and can customize it on a day-to-day basis.

It is recommended that the maintenance service be run daily for a period of at least four hours. By default, the maintenance service is run daily (including weekends) from 1 AM until 5 AM. This is a reasonable time frame for most users. If one wishes to change this, keep in mind the following points:

- The maintenance service can take up a significant amount of server resources, so it should be scheduled for periods when the load is expected to be light.
- The maintenance service should be allowed to run for at least 4 hours to ensure it completes. Larger systems may require more time. Monitor the maintenance activities and increase the duration of maintenance if regularly not completing within the allotted time.
- The maintenance interval should be run daily. Since the maintenance service handles final deletion of files, running the service less frequently can result in unnecessarily wasted memory.
- Ideally, the maintenance interval should take place after backups run. This is because some controls (for example,. item 4.3.2) prevent final deletion of files until after a backup has been performed. By running the backup first, the number of files that can be purged from the system is maximized.

If the maintenance interval is scheduled in such a way that it overlaps with a significant amount of user activity, users will likely notice a lag in service. If the maintenance duration is too short, or if it is run too infrequently or is not coordinated with the backups, then deleted files will continue to take up space on the server for longer than is necessary.

4.2.5 This Database can be Overwritten by a Restore

This controls whether the mailbox store can be overwritten by a backup. Stores that are overwritten by a backup will lose all information added after the backup was created.

This checkbox should be cleared. Doing this prevents accidental loss of data. The checkbox should only be selected immediately before a restore is to be made, and cleared again immediately afterwards.

Leaving this control selected opens up the possibility of data loss due to an accidental restore from backup

4.3 Deletion Settings

4.3.1 Do Not Permanently Delete Mailboxes Until Backed Up:
See item 4.2.3

4.3.2 Keep Deleted Items for (days)

This controls the minimum number of days that a deleted item (such as an email message) will be retained before it is purged from the system. When a user "deletes" a message, they are actually marking it for purging at a later point. The "deleted" message will, in fact, remain in the database until it is finally purged by the maintenance service. Until this happens, the message can still be recovered without resorting to backups. Other controls (such as item 4.3.1, "Do Not Permanently Delete Mailboxes Until Backed Up") can also influence how long a deleted item remains in memory. If multiple controls govern the purging of deleted messages, the message will not be deleted until the conditions specified by all the controls are met. For example, if deleted items are retained for 7 days and will not be purged until after a backup, if a deleted item is 8 days old but has not been backed up, it will remain until a backup occurs. This can be overridden on a per-user basis by item 12.1.1.

It is recommended that deleted messages be retained for 7 days before being purged. This strikes a balance between the desire to be able to recover deleted messages within a reasonable amount of time without resorting to backups, while at the same time reducing the amount of storage being consumed by deleted messages. Administrators may wish to increase or decrease this number if they find themselves frequently needing to recover deleted messages or if server storage becomes tight, respectively.

If the number of days a deleted message is retained is too small users wishing to retrieve deleted messages will need to have those messages retrieved from backups more often, which is time consuming both for the user and the administrator. On the other hand, if messages are retained too long, storage space will end up being wasted to hold unwanted mail. Administrators should monitor these factors and may wish to adjust the setting accordingly.

4.3.3 Keep Deleted Mailboxes for (days)

This controls the minimum number of days that a deleted mailbox will be retained before it is purged from the system. When a user "deletes" a mailbox, they are actually marking it for purging at a later point. The "deleted" mailbox will, in fact, remain in memory until it is finally purged by the maintenance service. Until this happens, the mailbox can still be recovered without resorting to backups. Other controls (such as item 4.3.1, "Do Not Permanently Delete Mailboxes Until Backed Up") can also influence how long a deleted mailbox remains in memory. If multiple controls govern the purging of deleted mailboxes, the mailbox will not be deleted until the conditions specified by all the controls are met. For example, if deleted mailboxes are retained for 30 days and will not be purged until after a backup, if a deleted mailbox is 31 days old but has not been backed up, it will remain until a backup occurs. This can be overridden on a per-user basis by item 12.1.1.

It is recommended that deleted mailboxes be retained for 30 days before being purged. This gives a large amount of flexibility to easily restore a user's mailbox. At the same time, since the messages in the mailbox will be deleted according to the timing specified by item 4.3.2. "Keep Deleted Items for (days)", the deleted mailbox will take up relatively little memory on the server.

If the mailbox is purged too quickly, the administrator will need to spend more time reconstructing the mailbox if a need for it arises again. If the number is too high, then the server will end up using storage space for an unused mailbox longer than necessary.

4.4 Display

4.4.1 Mailbox Store Policies

Mailbox Store Policies (created and stored under the Exchange System Manager → Administrative Groups → [administrative group] → System Policies container) can be used to control the settings configured in the Mailbox Store container's "General", "Database", "Limits", and "Full-Text Indexing" tabs. A Mailbox Store Policy can be applied to multiple mailbox stores in this administrative group, obviating the need to configure the stores individually and ensuring a uniform configuration across the various stores. If multiple mailbox stores are used, and if these mailbox stores are intended to have the same configuration (at least on the listed panels) then the use of Mailbox Store Policies is recommended both to simplify administration and to ensure configuration consistency.

This screen simply lists the Mailbox Store Policies that have been applied to this mailbox store. Use it to verify that a policy object has been successfully applied to the store and as a reminder of which tabs in the control panel are governed by policy objects. The policy itself cannot be manipulated through this interface.

4.5 Logging

4.5.1 Archive All Messages Sent or Received by Mailboxes

This controls whether messages that are received by or sent from this mailbox store should be archived. When the checkbox is selected you must also select a user, distribution list, contact, or public folder to whom all messages will be copied. This feature is more correctly called "Journaling" and is used to provide a "paper trail" of all correspondence that passes through the server. For more information on journaling, consult the article at http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/journaling.mspx.

This feature should be selected for accountability purposes. However, use of this feature also raises potential legal and privacy concerns. Organizational policy should be consulted before a final decision is made regarding this control. In addition, archiving all messages sent or received by mailboxes poses a significant burden on network traffic and memory usage on the journal machine. Journaled messages should always be stored on a separate dedicated journaling server.

If this control remains cleared, users will control the only copy (technically the only pointer to the copy) of this message on this mail store. While controls can be set to prevent messages from being permanently deleted until after they have been backed up (item 4.2.3), this feature provides another level of preservation for messages.

4.6 Permissions

4.6.1 Mailbox Store Permissions

This ACL controls the various rights to this mailbox store. Various rights include viewing the status of the store as well as changing its settings. Note that the "Special permissions" right will appear to be granted to most users and groups including the "Everyone" group. This right is required by Exchange, does not represent a security risk, and should not be changed.

Carefully scrutinize the permissions associated with each mailbox store and only include the most restrictive set of permissions needed. Due to the wide variety of possible configurations, no single recommendation for these settings can be given. Administrators must identify and grant required permissions on a store-by-store basis. That said, in general, the default configuration will suffice in many enterprises. The settings should only be changed with a solid understanding of the consequences.

4.7 Storage Limits

4.7.1 Storage Limits

These control the maximum sizes of a user's mailbox and the system's response if these limits are exceeded. There are three controls, each of which is intended to represent a different level of response depending on the degree to which the quota has been exceeded. The first control, meant to represent the first level of response, sends an email warning message to the user alerting them that they have exceeded their mailbox quota. At the second level, in addition to the warning message, the user will be prevented from sending email, although they will still be able to receive messages. At the third level, a warning message is sent and no further messages may be sent or received by the user. Note that these controls may be overridden on a per-user basis using the Active Directory Users and Computers interface (item 12.5.1). Each control may be selected or cleared. If cleared, the associated level or response will not occur. If all three controls are cleared the user's mailbox quota will be effectively unlimited. If the control is selected, a text box on the right is used to specify the mailbox size that must be exceeded for the associated response to occur.

It is recommended to select all three responses. If no responses are enabled then there is nothing to prevent excessive mail use by users and using all three levels results in a measured response to quota violations that provides users with a reasonable opportunity to respond before mail is lost. The specific limits at which each of the responses are triggered are completely a function of operational considerations (such as available disk space and expected email sizes) so we offer no recommendations on these settings. However, it is suggested that the differences between the response levels be at least as much as the maximum message size as controlled by SMTPVS (item 11.3.3). This prevents a single message from causing response levels to be skipped.

If no limits are applied to a user's mailbox, the mailbox size is effectively unlimited. This can result in excessive memory usage either due to users failing to expunge unneeded messages or due to mail bombing activities which fill the hard disk. If limits are enabled, at worst a mail bomber may be able to disable send and receive functionality for individual users, but they should be unable to affect the mailbox store as a whole. One item to note, if a user account becomes unused while exceeding quota, then warning messages will continue to be sent to the account, consuming additional space. For this reason, it is recommended that inactive accounts that cannot be deleted outright be brought below quota or be given custom quotas (using Active Directory Users and Computers – item 12.5.1) so that warning messages are not generated.

4.8 Mounting Stores

4.8.1 Do Not Mount This Store at Start-up

This controls whether this Mailbox Store should be mounted when Exchange starts. Stores are usually only unmounted when manual maintenance is being performed on them. When a store is unmounted, its contents are inaccessible to other users. Note that selecting this checkbox does not cause the store to become unmounted immediately – it only means that the store will not be re-mounted the next time Exchange starts.

This control should be cleared for general use. Doing this ensures that the store is mounted when Exchange starts and thus is accessible to users. If, however, conditions require that the store be unmounted (for example, maintenance), then this checkbox should be selected so that, should Exchange restart before maintenance is completed, it will not be inadvertently mounted in a bad state. Once the store is ready to mount again, the checkbox should be cleared so that the store will be remounted on boot as well.

If this checkbox is selected (the store will not be mounted on start-up) then the store will become unavailable the next time Exchange starts. If, however, maintenance is being performed that requires the store to be unmounted and this checkbox is *not* selected, then if the Exchange server reboots, it will attempt to mount the store automatically. If the store is in a bad state as a result of incomplete maintenance activities this may cause Exchange to fail to start correctly. As such, it is important that this control be set appropriately.

5. NNTP

5.1 Admin Post Requests

5.1.1 Allow Control Messages

Control messages correspond to three newsgroup folders: control.cancel, control.newgroup, and control.rmgroup. Posts to these newsgroups represent administrative requests. Specifically, posts to control.cancel represent requests to delete a specific post or article, posts to control.newgroup represent requests to create a new newsgroup, and posts to control.rmgroup represent requests to delete a newsgroup. If Allow Control Messages is selected, the requests that are sent to these newsgroups will be automatically processed by the NNTP Virtual Server. If this feature is cleared, the messages must be manually handled by an administrator.

Disable this feature. This forces all newsgroup creation and deletion to be explicitly performed by an Exchange administrator. This way, the ability to create and delete newsgroups remains the exclusive ability of administrators, instead of granted to anyone who can post to the special control newsgroups.

If this feature is enabled, the ability to post to these newsgroups implies the ability to create and delete newsgroups. Special care should be taken that only specifically allowed users and computers are allowed to post to these newsgroups (item 5.5.1). Moreover, the newsgroups should be moderated (item 5.4.4) so that messages will only appear in the control newsgroups after they have been reviewed by an administrator. If both of these actions are taken, much of the risk of this feature can be mitigated.

5.2 Authentication

5.2.1 Authentication Method Used to Access NNTP Virtual Server

This controls the authentication methods that are enabled for NNTP as well as the encryption that will be applied to the authentication. Options include Anonymous (no authentication), Basic (username and password) and Integrated Windows Authentication (also known as NTLM). In addition, the panel can require that SSL be used to encrypt the authentication, and can also request that users use client based certificate authentication. Specifically, if “Enable SSL client authentication” is selected, the communication between the client and server will be encrypted

using SSL. If “Require SSL client authentication” is set, then clients that support SSL certificate authentication will be required to use it. Note that, if the client does not support SSL certificate authentication, other authentication methods will still be accepted. As such, the name of this control is somewhat misleading. Finally, the “Enable client certificate mapping” control allows Exchange to associate certain client certificates with Windows identities. Note that no known clients currently support SSL certificate authentication, so enabling the latter two features will likely have no operational effect.

It is recommended that all features in this panel *except* “Allow anonymous” be selected and enabled. Allowing Anonymous authentication precludes the need of any individual to authenticate to the NNTP virtual server. Under most circumstances this would be undesirable. By default both Basic authentication and Integrated Windows Authentication (also known as, NTLM) are enabled and should remain enabled to ensure availability. Finally, “Enable SSL client authentication” should be enabled. This is particularly important since otherwise the username and password during basic authentication will be transmitted in clear text. “Require SSL client authentication” and “Enable client certificate mapping” may be enabled if desired but since no clients support these features, they are unlikely to be utilized.

If anonymous access is allowed then no restrictions will be placed on access to the NNTP virtual server and will lead to unrestricted access to posts. If Basic authentication or Integrated Windows Authentication are not enabled then some clients may not be able to authenticate to the NNTP server if they do not support the remaining authentication mechanism.

5.2.2 Certificate Wizard

This button starts a wizard to install a certificate to be used by this NNTP Virtual Server. Server certificates are required for many security features in Exchange, and without them the server cannot engage in many forms of secure communication. The wizard can guide the administrator through the process of requesting a new certificate or of importing an existing certificate. Certificates must be manually installed on each virtual server. This means that installing a certificate on one NNTP Virtual Server does not give other NNTP Virtual Servers (or the virtual server of any other protocol) access to this certificate. However, once a certificate is installed on one virtual server, any other virtual server (regardless of protocol used) may easily be configured to use this certificate by selecting “Assign an existing certificate” in the first page of the Wizard.

Install certificates on this virtual server. Without it, many other recommendations in this document concerning secure communication will be impossible. For highest security assurance, each virtual server should have its own certificate that it does not share with other servers. This reduces the damage due to server compromises and provides per-server identification.

Failure to implement this recommendation makes it virtually impossible to secure Exchange's communications. Use of any virtual server that has not been given a certificate should be considered a highly insecure action.

5.3 Connections
5.3.1 Require Secure Channel (SSL) and Require 128-bit encryption to NNTP Virtual Directory This controls whether client machines should be forced to use secure channels to communicate with this virtual server. If this feature is enabled, clients will only be able to download articles from the server if they are capable of supporting secure communication with the server. If secure channels are required, the server can also require the channel to be strongly secure by requiring 128 bit encryption. Require secure channels and 128-bit encryption. The use of secure communication prevents eavesdroppers from reading or modifying communications between servers and clients. If a particular client application in your organization does not support secure communication, this feature will need to be disabled or that client will be unable to connect to the server. However, it is strongly advised that any such clients be upgraded. Failure to enable this feature gives attackers more opportunities to read and modify communication between the client and server. Requiring secure channels but not requiring 128-bit encryption will still protect most messages from eavesdropping and modification, but makes it slightly more likely that a determined and well equipped attacker may be able to overcome the channel security.
5.3.2 Limit Number of Connections This controls the number of simultaneous connections that are permitted to this virtual server. If the checkbox is cleared then no limits are placed on the number of connections. If the checkbox is selected, the number in the text field specifies the maximum number of simultaneous connections that will be allowed. It is recommended that this value be set to 5000. This number should be large enough for most environments but small enough to reduce the risk that network performance will be degraded by a large number of connections. This said, some environments may have different requirements and may operate better with a different value. For example, if the enterprise suffers from a slower network connection then reducing the number of permitted connections may help reduce network congestion. On the other hand, a large enterprise with a correspondingly high bandwidth may wish to increase this number. Failure to limit the number of NNTP connections, or setting a limit that is too high may result in excessive network congestion and prevent users from accessing Exchange's services. On the other hand, a number that is too low may result in legitimate users being denied access to the NNTP virtual server because the maximum number of connections has already been reached. The recommendation provided here should be sufficient for most environments, but administrators should observe usage and modify the value accordingly.

5.3.3 Connection Time-out (Minutes)

This controls the number of minutes an idle connection will be maintained by the server before it is dropped. This can be used to help limit the number of simultaneous connections the server must support. The speed with which idle connections are dropped is particularly important if the number of simultaneous connections permitted (using item 5.3.2) is low, since idle connections may take up a large number of the available connections.

It is recommended that this value be set to 10 minutes. This provides users with a reasonable window in which to resume activities without maintaining idle connections for excessive intervals. Administrators should monitor usage and modify this value as necessary. If users are frequently having their sessions cut off unexpectedly, the value may need to be increased or the company may wish to invest in a speed reading class. If administrators find that they are frequently hitting the limit on the number of simultaneous connections provided in item 5.3.2 (but they feel this connection limit is reasonable for their environment) then the number of minutes until timeout should be reduced. That said, it is inadvisable to reduce the timeout much below 10 minutes as this will likely result in interruptions to users.

If connections time out after too few minutes then user sessions may be terminated unexpectedly. Idle periods are not uncommon with NNTP since users are likely to spend time reading a message before contacting the server again for another message. As such, idle periods of several minutes should be considered common. On the other hand, if the timeout period is too long, large number of unused connections may end up being maintained for an unnecessarily long period of time. This is unlikely to affect resources, since idle connections take up little processing power or bandwidth, but may result in the number of open connections hitting the limit prescribed by item 5.3.2, which could prevent new connections from being established.

5.4 Delivery Restrictions

5.4.1 Grant or Deny Access

This controls which IP addresses are allowed to connect to this virtual server. The control can be set to either allow all computers to connect except for a specified few, or to deny all computers except for a specified few. In addition to individual IP addresses, subnet masks may be used to specify groups of addresses. Note that if your network uses DHCP to dynamically assign IP addresses to client computers it will not be possible to reliably specify one computer out of the group serviced by a given DHCP server since that computer's IP address may change regularly. (If all computers serviced by a DHCP server should have the same treatment then there is no problem since one or more subnet masks can then be used to cover the DHCP server's entire address pool, ensuring that regardless of the address a computer is assigned, it will have the same treatment in this panel.)

Select “Only the list below” so that the administrator must explicitly specify which addresses can connect to the NNTP Virtual Server. This significantly reduces the change of unauthorized connections to the server and helps to further restrict network connectivity. It may not be possible to know the IP addresses of all possible clients, but if your system uses a VPN, the IP assigned by the VPN can be used here instead. This scenario is recommended if external access is required.

If “All except the list below” must be selected, administrators should monitor connectivity to the NNTP virtual server to ensure that no suspicious connections are being made.

5.4.2 Feed: Limit Post Size

Feed posting is the process by which one NNTP virtual server's newsgroups can be used as the source of postings to another NNTP virtual server's newsgroup. If “Allow feed posting” is selected on this server, than articles posted to this server can be propagated to additional NNTP virtual servers. This feature controls the maximum size of articles that can be downloaded via feed posting. If this feature is enabled, individual articles that exceed the size limit will not be forwarded to the destination server. Note that this checkbox will only be available if feed posting is enabled.

If feed posting is used, enable this feature and set the post size to at most 1500 KB. This can be used to ensure that network traffic to the NNTP virtual server from other NNTP virtual servers that are feeding off this server's posts do not consume excessive resources. If feed posting is not needed then there is no need to configure this control.

Failure to restrict the maximum post size can result in excessive demands on the network and server. If feed posting is not enabled then the risk is removed.

5.4.3 Feed: Limit Connection Size

Feed posting is the process by which one NNTP virtual server's newsgroups can be used as the source of postings to another NNTP virtual server's newsgroup. If “Allow feed posting” is selected, articles posted to this server can be propagated to additional NNTP virtual servers. This feature controls the maximum size session size to be used when downloading to an NNTP virtual server feeding from this server. If this feature is enabled, the feeding server will need to open additional sessions if it needs to download more data than the listed limit. Note that this checkbox will only be available if feed posting is enabled.

If feed posting is used, enable this feature and set the connection size to at most 40 MB. This can be used to ensure that network traffic to the NNTP virtual server from other NNTP virtual servers that are feeding off this server's posts do not consume excessive resources. If feed posting is not needed then there is no need to configure this control.

Failure to restrict the maximum post size can result in excessive demands on the network and server. If feed posting is not enabled then the risk is removed.

5.4.4 Moderated

This controls whether all messages to this particular newsgroup should be first directed to a specified user. Only after the messages have been reviewed will they be posted to the newsgroup for others to read. If this checkbox is selected, the email address of the moderator will need to be provided in the appropriate field below the checkbox. Care should be taken that the moderator's address is correct since, if messages cannot be delivered to the specified address, it will be impossible to post to the newsgroup.

While it is more secure to moderate, such a recommendation may be impractical due to the amount of traffic a newsgroup may have. It also assumes that the manpower is available to support such a feature. Having a moderator ensures that postings remain on topic, that the newsgroup does not end up reposting spam messages, and helps to ensure that sensitive or inappropriate material is not placed in the newsgroup. The moderator will need to be able to respond promptly to postings so appropriate articles can appear in the newsgroups in a timely manner. This feature is especially important for the three control message newsgroups discussed in item 5.1.1. In general, however, enterprises will need to evaluate which newsgroups require moderation and balance this against their ability to provide personnel to perform this task. As such, no recommendation is provided here.

If the moderator checkbox is not selected, all posts to this newsgroup will be immediately displayed for all newsgroup readers. For this reason, it is important that access to read and

post to the newsgroup be tightly controlled (items 5.5.1 & 7.6.1). Moreover, a newsgroup's administrator should regularly read all new posts to limit the amount of time inappropriate material remains in the newsgroup. Administrators will also need to be available to quickly respond to user complaints regarding the posting of inappropriate or offensive material since failure to respond promptly in such situations can result in legal liability.

5.4.5 Allow Posting

This controls whether clients are allowed to post to this newsgroup. It can only be used to turn posting on or off - it cannot be used to restrict access to a specific list of users. (To do this, use the Windows file security of the appropriate file - item 5.5.1 or use the Public Folder Store client permissions –item 8.3.1, depending on how the virtual directory is stored.) Selecting the checkbox allows users to post to this virtual directory. Otherwise, no users are permitted to post to the directory.

Although it would be more secure to disallow posting, doing so breaks the purpose of a newsgroup. As such, we provide no recommendation for this control. Items 5.5.1 and 8.3.1 can help to limit access to a newsgroup. If a newsgroup is no longer active (it has been frozen or is no longer relevant) disallowing posting can help protect it against further modifications.

If posting is allowed, the newsgroup should be monitored to ensure that the newsgroup is being used appropriately.

5.4.6 Restrict Newsgroup Visibility

This controls whether users will be allowed to see the names of newsgroups to which they do not have read access. Read and write access for individual users is controlled using NTFS file security controls (item 5.5.1) or Public Folder Store client permissions (item 8.3.1). If this checkbox is not selected, the newsgroup's name will be part of a list that the Exchange server sends out to all users, regardless of their ability to read the newsgroup's contents. If selected, the user's access to the directory is selected before the newsgroup's name is displayed.

Select this feature. By doing this, one limits the amount of information about the system configuration that is unnecessarily available. By hiding the existence of the folder, added protection is provided against those who would wish to read or modify the folder's contents.

Failure to implement this feature means that all users will know of the folder's existence. This information can then be used to plan attacks to read or modify the folder contents.

5.4.7 Allow Client Posting: Limit Post Size (kb)

If “Allow client posting” is selected on this virtual server, clients will be allowed to post to any of the contained newsgroups to which they are not otherwise denied access. Clients might be denied access by the directory’s permissions (items 5.5.1 and 8.3.1) or if the directory does not allow posting (item 5.4.5). If client posting is allowed, this item controls the maximum size of messages that can be posted to this virtual server. Preventing excessively large files from being posted not only helps slow server memory consumption, but also ensures that clients will not get bogged down trying to download large files to their local computer. If the checkbox next to this control is cleared, there is no limit on the size of posts. If it is selected, the text field on the right specifies the maximum size of posts.

It is recommended that message sizes be limited to 100K. This should be sufficient for most standard uses. If organizations expect for larger files to be posted (such as images, movies, sound clips, or other such files), the maximum post size may need to be increased in order to accommodate them.

If no limits are placed on message sizes, or if the limit is too large, then large messages may end up overwhelming the server and/or client abilities to download and read the posts. If, however, the limit is too low, users may be unable to post legitimate messages to the server. If client posting is not enabled there is no risk from this control’s setting.

5.4.8 Allow Client Posting: Limit Connection Size (mb)

If “Allow client posting” is selected on this virtual server, clients will be allowed to post to any of the contained newsgroups to which they are not otherwise denied access. Clients might be denied access by the directory’s permissions (items 5.5.1 & 8.3.1) or if the directory does not allow posting (item 5.4.5). If client posting is allowed, this item controls the maximum size of connections to the virtual server. This can be used to reduce the number of messages that can be posted in a single connection. If more information needs to be posted, additional sessions will need to be opened. This can make it slightly more difficult to bombard a newsgroup with a large number of posts. If the checkbox for this control is cleared, no limit will be applied to session sizes. If the checkbox is selected, the text box to the right specifies the maximum session size permitted in MB.

It is recommended that connection size be limited to 1MB. This should be sufficient for most enterprises. If a large number of legitimate sessions are being rejected, or if too many sessions must be created during normal usage, the session size may need to be increased. Likewise, if the message size limit (item 5.4.7) is increased it may be advisable to increase the session size. At the very least, the session size should be larger than the maximum message size, since no message larger than the session limit can be delivered.

If the session size is unlimited or too large, the server may be inundated by exceptionally large communication sessions. If it is too small, however, at best it may force servers to break communications into a large number of small sessions, and at worst it may completely prevent large but legitimate messages from being delivered. If client posting is not enabled there is no risk from this control’s setting.

5.5 Directory Access

5.5.1 Physical Location Security Panel

NNTP virtual directories may be stored either in the public folder store or in their own directory on the hard disk. Exchange controls user access to directories in the folder store using item 8.3.1. To enforce per-user control of NNTP directories that are stored separately on the disk, security must be applied to the NTFS security tab for the virtual directory folder within the Windows file structure. By limiting the users who are able to read or write the folder within the Windows file system, access is also limited via the NNTP protocol. Virtual directories usually appear under \Inetpub\nntpfile\root\, although they can be configured to appear anywhere, and have the same name as the virtual directory. (The virtual directory properties control will list the directory's location.) Security only needs to be applied to the directory itself, not to any of the contents.

Restrict access to the virtual directory. However, it is advisable that someone with administrative responsibilities always have full access to the virtual directories so they can be administered.

Failure to limit access here means that all users will have the same level of access to the folder. While this may make sense in some cases (such as for public newsgroups), it is inadvisable for newsgroups that may contain sensitive or private information.

Note that this only can be used to control access to newsgroups that are stored separately instead of on the Public Folder Store. To control access to newsgroups on the Public Folder Store, use item 8.3.1.

5.6 Logging

5.6.1 Enable Logging

This controls the creation and format of log files used to monitor transactions with this virtual server. By default, these files will be stored in \WINNT\SYSTEM32\LOGFILES\NntpSvcx (where x is a number used to distinguish between virtual servers in this organization). The drop-down menu is used to select the format of the log file. The properties button next to this dropdown displays configuration information specific to the type of log format selected, but usually has some control to indicate the log rotation schedule (that is, how often the old log file should be closed and a new log file should be started). The log files contain, among other things, the times that users connect, what articles they read, and from where they are connecting. Note that this only enables logging at the virtual server level. Logging must also be enabled at the virtual directory level (item 5.6.2) for any log data to actually be recorded.

Enable logging. Doing so records attempts made to connect to this virtual server. In the case of an attack on the NNTP server, these logs could contain useful details regarding the time and nature of the attack. Due to the size of log files, the files should be regularly copied to external storage and deleted from the server to conserve memory. Log files should be retained for at least one month. The format of the log files is largely a matter of preference for the administrator. Due to the degree to which user activity is recorded by the logs (including which articles users read) the organization's privacy policy may need to be consulted before this feature is enabled. If enabled, care should be taken that log files are not accessible to non-administrators, both for privacy reasons, and because the log files can provide attackers with information they can use to plan their attacks.

Failure to enable this feature means that the administrator will not have access to information regarding NNTP connections should the NNTP servers come under attack.

It is required that all log files be written to separate partitions from the partitions used by the Exchange Stores. This drive should be dedicated to log files as they can be quite large and fill up quickly. Exchange will dismount its stores if it detects that it has run out of disk space,

resulting in a complete loss of Exchange services. To minimize the chance of this happening, log files should write to some other partition so that if the logs fill this partition it will not result in the failure of Exchange.

5.6.2 Log Access

This feature enables transaction logging for this virtual directory. This control only takes effect if logging is enabled for the virtual server (item 5.6.1) and log entries will appear in the virtual server log file.

Enable this feature so that transactions with this virtual directory are recorded. The transaction logs can be used to track down abuse of the NNTP service as well as to help provide details about attacks that were directed through this virtual directory.

Failure to enable this feature will mean that no transactions with this virtual directory will be recorded. The virtual server can still record transaction information for other virtual directories (assuming that logging is enabled on them) regardless of this setting. However, disabling logging of this virtual directory limits the amount of information available to respond to emergencies.

It is required that all log files be written to separate partitions from the partitions used by the Exchange Stores. Exchange will dismount its stores if it detects that it has run out of disk space, resulting in a complete loss of Exchange services. To minimize the chance of this happening, log files should write to some other partition so that if the logs fill this partition it will not result in the failure of Exchange.

6. POP3

6.1 Authentication

6.1.1 Authentication Method used to Access POP3 Virtual Directory

This controls the form of authentication and encryption used by clients attempting to connect to this virtual server. Possible settings are “Basic authentication” and “Simple Authentication and Security Layer”. If Basic authentication is selected, an additional checkbox allows the server to require that SSL/TLS encryption be used to protect the password (which would otherwise be sent in clear text). Simple Authentication and Security Layer is the Microsoft NTLM protocol.

Select the “Basic authentication” and “Require SSL/TLS” checkboxes. NTLM can be secured, but it is a negotiated protocol and can lead to variable levels of security. The use of SSL/TLS not only protects the username and password during authentication, but encrypts the mail messages as they are being transmitted, preventing eavesdroppers from reading or modifying messages.

Failure to implement SSL/TLS with Basic authentication is extremely insecure as it results in the user's identity and password being sent over the network in clear text. If this happens, network sniffers can easily collect these credentials. The use of NTLM (Simple Authentication and Security Layer checkbox), while it can protect the username and password during authentication, it does not provide encryption of message bodies, potentially allowing message content to be sniffed over the wire. Moreover, NTLM negotiates the details of how it secures the authentication and some possible selections are insecure. While it is possible to configure NTLM server to only use the more secure options, Exchange does not provide this capability. (To only permit the most secure form of the protocol, NTLM v2, set the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\LMCompatibility registry key to 5. See Microsoft Knowledge Base Article 239869.)

6.1.2 Certificate Wizard

This button starts a wizard to install a certificate to be used by this POP3 Virtual Server. Server certificates are required for many security features in Exchange, and without them the server cannot engage in many forms of secure communication. The wizard can guide the administrator through the process of requesting a new certificate or of importing an existing certificate. Certificates must be manually installed on each virtual server. This means that installing a certificate on one POP3 Virtual Server does not give other POP3 (or any other protocol) virtual servers access to this certificate. However, once a certificate is installed on one virtual server, any other virtual server (regardless of protocol used) may easily be configured to use this certificate by selecting “Assign an existing certificate” in the first page of the Wizard.

Install certificates on this virtual server. Without it, many other recommendations in this document concerning secure communication will be impossible. For highest security assurance, each virtual server should have its own certificate that it does not share with other servers. This reduces the damage due to server compromises and provides per-server identification.

Failure to implement this recommendation makes it virtually impossible to secure Exchange's communications. Use of any virtual server that has not been given a certificate should be considered a highly insecure action.

6.2 Connections

6.2.1 Require Secure Channel (SSL) and Require 128-bit encryption to POP3 Virtual Server

This controls whether client machines should be forced to use secure channels to communicate with this virtual server. If this feature is enabled, clients will only be able to download mail from the server if they are capable of supporting secure communication with the server. If secure channels are required, the server can also require the channel to be strongly secure by requiring 128 bit encryption.

Require secure channels and 128-bit encryption. The use of secure communication prevents eavesdroppers from reading or modifying communications between servers and clients. While sensitive message bodies can be encrypted by the sender at the client, securing the communication between the server and the recipient provides an added level of security while also protecting sender and recipient information that cannot be encrypted by the sender.

If a particular client application in your organization does not support secure communication, this feature will need to be disabled or that client will be unable to connect to the server. However, it is strongly advised that any such clients be upgraded. Failure to enable this feature gives attackers more opportunities to read and modify communication between the client and server. Requiring secure channels but not requiring 128-bit encryption will still protect most messages from eavesdropping and modification, but makes it slightly more likely that a determined and well equipped attacker may be able to overcome the channel security.

6.2.2 Use SSL Connections

This controls whether responses to meeting requests will use an encrypted channel when responding to servers. Meeting request emails contain links to an Outlook Web Access page that allows users to access most of the features of standard Outlook meeting requests over the web. The “Use SSL connections” checkbox determines whether the download and subsequent server interaction via the downloaded form uses a secure channel.

Select the “Use SSL connections” checkbox. This protects user interactions with their calendar from being observed or modified by an eavesdropper.

Failure to select this feature means that the use of OWA to handle meeting requests may be observed or modified by a third party.

Note that for this to work, the HTTP “Exchange” Virtual Directory must be present and configured to use secure channels. (IIS Manager → [server] → Web Sites → Exchange → Properties → Directory Security Tab → Secure Communications Edit button → Require secure channel (SSL) checkbox).

6.2.3 TCP Port/SSL Port

This controls the ports that the POP Virtual Server binds to for regular and secure communications, respectively. The standard port for regular POP connections is 110, while the standard port for secure POP is 995. If ports other than these are to be used, clients will need to be explicitly configured to use the non-standard port.

Use standard ports for both protocols (that is, the default). Traffic on these ports should be carefully monitored for signs of suspicious activity.

Changing the port to a non-standard value can provide some limited protection against automated attacks since these attacks will not connect to the correct port, and therefore be nullified. However, making this modification introduces a large amount of complexity for the system administrator: If clients are to connect to the server, they will all need to be configured to use the selected non-standard ports for their communication. In addition, if at some later point, additional network services are added to the Exchange server, care must be taken that the ports used by these services do not conflict with the non-standard ports chosen here. Finally, if clients outside the firewall or proxy server are to be allowed to connect to the server, or if automated network monitoring tools are employed, these services will all need to be configured to use the non-standard ports for POP traffic. Even if ports are changed, however, a determined attacker may still be able to determine which port is used for the POP protocol by performing a comprehensive port scan, although doing so should be detectable by most network monitoring tools. Since changing the port introduces a large amount of complexity for a relatively small gain, the standard ports should be used.

6.2.4 Limit Number of Connections

This controls the maximum number of simultaneous connections allowed to the POP server. This can be used to throttle the use of the POP service if it begins to monopolize system resources. If the checkbox is cleared, no limits are applied to the number of connections. If the checkbox is selected, the text field on the right is used to specify the maximum number of connections.

The setting of this control is highly dependent on environment. In some cases, applying a limit to the number of POP connections is a useful way to control the network resources that are consumed by this protocol. This is especially likely if POP is not the primary role of this server. However, in other cases, applying a limit may be detrimental to the enterprise's mission because it may end up prohibiting legitimate POP use. This is particularly likely if the server is primarily intended to service POP requests and network resources are unlikely to be a limiting factor. Because of the range of scenarios that would affect the setting of this control, no recommendation is given here.

If no limit is provided, the administrator should closely monitor POP use to ensure that POP activity does not come at the expense of other protocols. If the administrator notices this happening, limiting the number of POP connections may help remediate the problem. (This symptom might, however simply indicate that the server's network resources are insufficient for its use.) Likewise, if a limit is applied, administrators should be careful that this limit is not so low that it needlessly limits legitimate POP connections.

6.2.5 Connection Time-out (Minutes)

This controls the number of minutes an idle connection will be maintained by the server before it is dropped. This can be used to help limit the number of simultaneous connections the server must support. The speed with which idle connections are dropped is particularly important if the number of simultaneous connections permitted (using item 6.2.4) is low, since idle connections may take up a large number of the available connections.

It is recommended that this value be set to 10 minutes. This provides users with a reasonable window in which to resume activities without maintaining idle connections for excessive intervals. Administrators should monitor usage and modify this value as necessary. If users are frequently having their sessions cut off unexpectedly, the value may need to be increased. If administrators find that they are frequently hitting the limit on the number of simultaneous connections provided in item 6.2.4 (but they feel this connection limit is reasonable for their environment) then the number of minutes until timeout should be reduced. That said, it is inadvisable to reduce the timeout much below 10 minutes as this will likely result in interruptions to users.

If connections time out after too few minutes then user sessions may be terminated unexpectedly. On the other hand, if the timeout period is too long, large number of unused connections may end up being maintained for an unnecessarily long period of time. This is unlikely to affect resources, since idle connections take up little processing power or bandwidth, but may result in the number of open connections hitting the limit prescribed by item 6.2.4, which could prevent new connections from being established.

6.3 Delivery Restrictions

6.3.1 Exclude or Limit Connections

This controls which IP addresses are allowed to connect to this virtual server to download messages. The control can be set to either allow all computers to connect except for a specified few, or to deny all computers except for a specified few. In addition to individual IP addresses, subnet masks may be used to specify groups of addresses. Note that if your network uses DHCP to dynamically assign IP addresses to client computers it will not be possible to reliably specify one computer out of the group serviced by a given DHCP server since that computer's IP address may change regularly. (If all computers serviced by a DHCP server should have the same treatment then there is no problem since one or more subnet masks can then be used to cover the DHCP server's entire address pool, ensuring that regardless of the address a computer is assigned, it will have the same treatment in this panel.)

Select “Only the list below” so that the administrator must explicitly specify which clients can connect to the POP Virtual Server. This significantly reduces the chance of unauthorized connections to the server and helps to further restrict network connectivity.

If “All except the list below” must be selected, administrators should monitor connectivity to the POP server to ensure that no suspicious connections are being made.

Note that, for proper operation, all back-end servers must be able to connect to all other back-end servers within their cluster as well as their corresponding front-end server. (That is, the IP address of these servers must be allowed in this panel). Likewise, all front-end servers must be able to connect to all legal clients as well as all back-end servers. (It may not be possible to know the IP addresses of all possible clients, but if your system uses a VPN, the IP assigned by the VPN can be used here instead. This scenario is recommended if external access is required.)

7. Public Folder Store

7.1 Backup/Restore

7.1.1 Time of Last Full Backup

This field lists the time of the most recent known full backup of this Public Folder Store. To be "known", the backup must either use the Windows Server Backup Wizard (Start → All Programs → Accessories → System Tools → Backup) or some other Windows/Exchange aware backup utility. Only Windows/Exchange aware backup utilities are recommended since other backup utilities may not be able to correctly restore the mailbox store. This field cannot be set - it is simply present to provide information. Review this field soon after each scheduled full backup to ensure that the backup completed successfully.

Full backups of the public folder store should occur at least on a weekly basis. More frequent full backups will simplify any data restoration that may be necessary. Public folder store backups should take place with or in addition to backups of the full server.

7.1.2 Time of Last Incremental Backup

This field lists the time of the most recent known incremental backup of this Public Folder Store. To be "known", the backup must either use the Windows Server Backup Wizard (Start → All Programs → Accessories → System Tools → Backup) or some other Windows/Exchange aware backup utility. Only Windows/Exchange aware backup utilities are recommended since other backup utilities may not be able to correctly restore the mailbox store. This field cannot be set - it is simply present to provide information. Review this field soon after each scheduled incremental backup to ensure that the backup completed successfully.

Incremental backups of the public folder store should occur at least on a daily basis. More frequent backups can help to minimize the amount of work that is lost in case of a disk failure, but may result in a significant load on the server. The frequency of incremental backups should balance the importance of the information on the store against the ability of the server to handle incremental backups simultaneously with its other duties.

Public folder store backups should take place with or in addition to backups of the full server.

7.1.3 Maintenance Interval

This controls when and for how long the Exchange maintenance service will run. The maintenance service performs a number of activities including final deletion of files (as governed by other policy controls) and checking folder storage limits against policy. The administrator is given a very fine degree of control of when the maintenance service runs and can customize it on a day-to-day basis.

By default, the maintenance service is run daily (including weekends) from 1 AM until 5 AM. This is a reasonable time frame for most users. If one wishes to change this, keep in mind the following points:

- The maintenance service can take up a significant amount of server resources, so it should be scheduled for periods when the load is expected to be light.
- It is suggested that the maintenance service be allowed to run for at least 4 hours to ensure it completes. Larger systems may require more time. Monitor the maintenance activities and increase the duration of maintenance if regularly not completing within the allotted time.
- The maintenance interval should be run daily. Since the maintenance service handles final deletion of files, running the service less frequently can result in unnecessarily wasted memory.
- Ideally, the maintenance interval should take place after backups run. This is because some controls (item 7.3.2) prevent final deletion of files until after a backup has been performed. By running the backup first, the number of files that can be purged from the system is maximized.

If the maintenance interval is scheduled in such a way that it overlaps with a significant amount of user activity, users will likely notice a lag in service. If the maintenance duration is too short, or if it is run too infrequently or is not coordinated with the backups, then deleted files will continue to take up space on the server for longer than is necessary.

7.1.4 This Database Can be Overwritten by a Restore

This controls whether the folder store can be overwritten by a backup. Stores that are overwritten by a backup will lose all information added after the backup was created.

This checkbox should be cleared. Doing this prevents accidental loss of data. The checkbox should only be selected immediately before a restore is to be made, and cleared again immediately afterwards.

Leaving this control selected opens up the possibility of data loss due to an accidental restore from backup.

7.1.5 Do Not Permanently Delete Items until Backed-Up

This control allows the administrator to prevent the permanent deletion of public folders and their contents until after they have been backed up (using Exchange aware backup utilities such as the Microsoft Backup wizard from Start → All Programs → Accessories → System Tools → Backup). Exchange handles folder deletion by keeping a count of the number of references to each individual folder on the folder store. When a user "deletes" a folder item they are actually just removing their reference to that item. Once all the references to an item have been deleted, that item is marked as being ready for final deletion. A maintenance process (scheduled by item 7.1.3) runs periodically and actually deletes all items that are marked for final deletion. It is only at this point that the item is actually deleted from the store and its memory made available. This checkbox ensures that at least one backup has been run on the folder store before the item actually disappears.

Enable this feature by selecting the checkbox since it significantly reduces the chance of accidental deletion of important information. It also means that all items written to recipients who have accounts on this store will reside in backups even after deletion.

The only potential drawback to this is if the backup facility is inoperable for an extended period of time, the size of the folder store may grow due to the fact that no items are actually being deleted. Even under these circumstances, however, it is probably more important to wait until the backup service is restored then risk permanently losing important information.

7.2 Connections

7.2.1 Clients Support S/MIME Signatures

This checkbox is set if the email clients serviced by this store support the use of S/MIME signatures. If selected, Exchange will use S/MIME signatures to verify the source of messages.

All clients in the enterprise should be updated to support S/MIME. If there is a valid reason that S/MIME cannot be used, other client-specific mechanisms for message signing should be used. However, S/MIME is one of the most commonly supported mechanisms so its use is most likely to be compatible with recipient mail clients.

7.3 Deletion Settings

7.3.1 Keep Deleted Items for (days)

This controls the minimum number of days that a deleted item will be retained before it is purged from the system. When a user "deletes" an item, they are actually marking it for purging at a later point. The "deleted" message will, in fact, remain in memory until it is finally purged by the maintenance service. Until this happens, the item can still be recovered without resorting to backups. Item 7.1.5 "Do Not Permanently Delete Items until Backed Up" can also influence how long a deleted item remains in memory. If multiple controls govern the purging of deleted items, the items will not be deleted until the conditions specified by all the controls are met. For example, if deleted items are retained for 7 days and will not be purged until after a backup, if a deleted item is 8 days old but has not been backed up, it will remain until a backup occurs. This control can be overridden on a per-folder basis using item 8.1.1.

It is recommended that deleted items be retained for 7 days before being purged. This strikes a balance between the desire to be able to recover deleted items within a reasonable amount of time without resorting to backups, while at the same time reducing the amount of storage being consumed by deleted items. Administrators may wish to increase or decrease this number if the find themselves frequently needing to recover deleted items or if server storage becomes tight, respectively.

If the number of days a deleted items is retained is too small, users wishing to retrieve deleted items will need to have those messages retrieved from backups more often, which is time consuming both for the user and the administrator. On the other hand, if items are retained too long, storage space will end up being wasted to hold unwanted material. Administrators should monitor these factors and may wish to adjust the setting accordingly.

7.3.2 Do Not Permanently Delete Items until Backed-Up: See item 7.1.5

7.3.3 Age Limit for All Folders in This Store (Days)

This controls the number of days that a message will remain in any of the folders in this store before it is automatically marked for deletion. Public folders may be used to contain things that are only valid for limited periods of time, such as announcements. This control can be used to automatically mark such messages for deletion so administrators do not end up needing to review and delete the items manually. If the checkbox is not selected, messages will not be automatically marked for deletion regardless of their age. If the checkbox is selected, messages whose age exceeded the number of days specified by the field on the right will be marked for deletion. The setting of this control can be overridden on a per-folder basis by item 8.1.2, but note that item 8.1.2 can only change the number of days before a message is deleted but cannot disable automatic deletion entirely if it is enabled on the store. Note that this control only governs when the items become marked for deletion. Items such as 7.3.1 and 7.3.2 continue to govern the period between the time an item is marked for deletion and the time it is finally purged from the system.

It is recommended that age limits be disabled by leaving the checkbox cleared. This prevents items from being marked for deletion unintentionally. If a particular directory is intended to contain only short-term items, then use item 8.1.2 to enable age limits on that directory only. If age limits are set at the store, then every folder on the system will automatically delete items after some interval, which is likely to be undesirable. If this recommendation is followed, administrators (or some authorized individual) should periodically review the contents of all the folders in this store and delete unnecessary material to help reduce the folder size. This is especially important if the folder has a size limit (items 7.7.1 or 8.4.1).

If age limits are enabled, then administrators should take care that any folders on this store whose contents should not be deleted automatically override the default value with a suitably large expiration time using item 8.1.2 and that these settings are verified on a regular basis to avoid accidental changes. The administrator should also ensure that any new folders have correct age limits assigned.

7.4 Display

7.4.1 Public Folder Policies

Public Folder Store Policies (created and stored under the Exchange System Manager → Administrative Groups → [administrative group] → System Policies container) can be used to control the settings configured in the Public Folder Store container's “General”, “Database”, “Limits”, “Replication”, and “Full-Text Indexing” tabs. A Public Folder Store Policy can be applied to multiple public folder stores in this administrative group, obviating the need to configure the stores individually and ensuring a uniform configuration across the various stores. If multiple public folder stores are used, and if these public folder stores are intended to have the same configuration (at least on the listed panels) then the use of Public Folder Store Policies is recommended both to simplify administration and to ensure configuration consistency.

This screen simply lists the Public Folder Store Policies that have been applied to this public folder store. Use it to verify that a policy object has been successfully applied to the store and as a reminder of which tabs in the control panel are governed by policy objects. The policy itself cannot be manipulated through this interface.

7.5 Mounting Stores

7.5.1 Do Not Mount This Store at Start-up

This controls whether this Public Folder Store should be mounted when Exchange starts up. Stores are usually only unmounted when manual maintenance is being performed on them. When a store is unmounted, its contents are inaccessible to other users. Note that selecting this checkbox does not cause the store to become unmounted immediately – it only means that the store will not be re-mounted the next time Exchange starts.

This control should be cleared for general use. Doing this ensures that the store is mounted when Exchanges starts and thus is accessible to users. If, however, conditions require that the store be unmounted (for example, maintenance), then this checkbox should be selected so that, should Exchange restart before maintenance is completed, it will not be inadvertently mounted in a bad state. Once the store is ready to mount again, the checkbox should be cleared so that the store will be remounted on boot as well.

If this checkbox is selected (the store will not be mounted on start-up) then the store will become unavailable the next time Exchange starts. If, however, maintenance is being performed that requires that the store be unmounted and this checkbox is *not* selected, then if the Exchange server reboots, it will attempt to mount the store automatically. If the store is in a bad state as a result of incomplete maintenance activities this may cause Exchange to fail to start correctly. As such, it is important that this control be set appropriately.

7.6 Permissions

7.6.1 Public Folder Permissions

This ACL controls the various rights to this public folder store. Rights include viewing the status of the store as well as changing its settings. Note that the “Special permissions” right will appear to be granted to most users and groups including the “Everyone” group. This right is required by Exchange, does not represent a security risk, and should not be changed.

Carefully scrutinize the permissions associated with each public folder store and only include the most restrictive set of permissions needed. Due to the wide variety of uses of public folders and public folder stores, no single recommendation for these settings can be given. Administrators must identify and grant required permissions on a store-by-store basis. That said, in general, the default configuration will suffice in many enterprises. The settings should only be changed with a solid understanding of the consequences.

7.7 Storage Limits

7.7.1 Issue Warning at (kb)/Prohibit Post at (kb)

These control the maximum sizes of a public folder and the system's response if these limits are exceeded. There are two controls, each of which is intended to represent a different level of response depending on the degree to which the quota has been exceeded. The first control, meant to represent the first level of response, sends an email warning message to users with Owner or Folder Contact roles for this public folder alerting them that the folder has exceeded its quota. At the second level, posting is no longer allowed to the folder. (Note that the description of this particular control in the help window in Exchange is incorrect.) Each control may be selected or cleared. If cleared, the associated level or response will not occur. If both controls are cleared the user's mailbox quota will be effectively unlimited. If the control is selected, a text box on the right is used to specify the folder size that must be exceeded for the associated response to occur.

It is recommended that both responses be selected. If no responses are enabled then there is nothing to prevent excessive posting to this folder and using both levels results in a measured response to quota violations that provides folder administrators with a reasonable opportunity to respond before posting is prohibited. The specific limits at which each of the responses is triggered are completely a function of operational considerations (such as available disk space and expected post sizes) so we offer no recommendations on these settings. However, it is suggested that the differences between the response levels be at least as much as the maximum post size as specified in item 7.7.2. This prevents a single post from causing response levels to be skipped.

If no limits are applied to a public folder, the folder size is effectively unlimited. This can result in excessive memory usage either due to failing to expunge unneeded messages or due to mail bombing activities which fill the hard disk and possibly disable the folder store. If limits are enabled, at worst a mail bomber may be able to disable posting functionality for individual folder, but they should be unable to affect the folder store as a whole.

7.7.2 Maximum Item Size (kb)

This controls the maximum size of any single post to a folder on this store. It is intended to prevent overly large messages from being placed in the folder. If the checkbox is cleared, there is no limit to the maximum size of a single post (although the limits to the total size of the public folder as set in item 7.7.1 will be observed). If the checkbox is selected, then posts larger than the value given in the text box at the right will be rejected. This control can be overridden on a per folder basis by item 8.4.1.

It is recommended that a maximum post size be enforced (by selecting the checkbox) and the maximum post size be set to 10240 KB (10 MB). Depending on network and storage capacity, as well as expected use of the folders in this store, administrators may wish to modify this. As a guideline, set the limit to a value that would be reasonable for most new folders, and then override for special case folders that have different requirements.

If no limit is placed on the size of posts then a single post may be able to cause the folder's quota to be exceeded (as set by item 7.7.1) or, if no limit is set on the folder size, a series of such messages may consume all the memory available to the folder store. It is important, however, to also keep operational requirements in mind when setting this limit so large but reasonable files are not rejected.

8. Public Folders

8.1 Deletion Settings

8.1.1 Deletion Settings

This controls whether this public folder uses the default deletion settings as set on the public folder store (item 7.3.1) or if the default should be overridden. The checkmark next to “Use public store defaults” controls whether the public store settings are overridden. If this checkbox is selected, then the public folder store’s settings are applied to this folder and the other field in this group is disabled. If the checkbox is cleared, then the settings provided by the store are discarded and the remaining field in this group (“Keep deleted items for”) is used to control the length of time a deleted item will remain in the store before being purged. The latter field functions in the same way as the corresponding field in the public folder store (item 7.3.1). Please consult this control for recommendations and risks regarding its setting. Note that this control only allows one to change the number of days before a deleted post is purged (item 7.3.1). Please consult this control for recommendations and risks regarding this setting. There is no control to override the setting that controls whether a backup must be performed before deleted items are purged (item 7.3.2).

It is recommended that, in general, public folders use the setting prescribed by the public folder store. There may be cases where it is reasonable or necessary to change the settings of a particular folder, but these cases should be limited as much as possible.

If a folder overrides the settings of the public folder store then it needs to be administered independently. This means that, when operational or capacity requirements change, in addition to updating the settings on the store, settings of all the independently administered folders must be updated as well. This increases the chance that the folder may be overlooked and end up configured in a way that is no longer appropriate. While there are certainly circumstances where overriding the settings of the folder store is the correct course of action, doing so does require additional time and attention on the part of the administrator.

8.1.2 Age Limits

This controls whether this public folder uses the default age limits as set on the public folder store (item 7.3.3) or if the default should be overridden. The checkmark next to "Use public store defaults" controls whether the public store settings are overridden. If this checkbox is selected, then the public folder store's settings are applied to this folder and the other field in this group is disabled. If the checkbox is cleared, then the settings provided by the store are discarded and the remaining field in this group "Age limit for replicas" determines any age limits for folder items. The latter field functions the same way as the "Age Limit for All Folders in This Store" (item 7.3.3) in the public folder store. (This control refers to replicas since folders may have different replicas of the same item in which case the replicas are tracked independently.) Please consult this control for recommendations and risks regarding this setting. Note that this control only allows you to change the age limit of items in the folder but does not allow you to disable age limits on this folder. If age limits are set on the public folder store, but you do not want items to be automatically deleted the only option is to set the age limits for this folder to a very high number. Setting the field to the maximum allowed value of 24855 (or just over 65 years) is probably sufficient.

It is recommended that, in general, public folders use the setting prescribed by the public folder store. There may be cases where it is reasonable or necessary to change the settings of a particular folder, but these cases should be limited as much as possible.

If a folder overrides the settings of the public folder store then it needs to be administered independently. This means that, when operational or capacity requirements change, in addition to updating the settings on the store, settings of all the independently administered folders must be updated as well. This increases the chance that the folder may be overlooked and end up configured in a way that is no longer appropriate. While there are certainly circumstances where overriding the settings of the folder store is the correct course of action, doing so does require additional time and attention on the part of the administrator.

8.2 Delivery Restrictions

8.2.1 Sending Message Size

This controls whether this folder will use the default sending message size limit (set in item 1.2.4) or whether the folder will observe different limits. If "Use default limit" is selected, the folder will observe the sending message size limit set in the global policy in item 1.2.4. If "Maximum KB" is selected, the sending message size limit as set in the global settings will be overridden and the text field on the right side of the panel will be used to hold the folder's customized sending size limit. See item 1.2.4 for recommendations regarding message size limits and the risks of various settings.

Note that the default Internet Newsgroups public folder and other folders that are not mail enabled do not have this tab or any of the controls contained therein. As such, configuring these folders is neither necessary nor possible.

It is recommended that, in general, folders use the setting prescribed in the global settings. There may be cases where it is reasonable or necessary to change the settings of a particular folder, but these cases should be limited as much as possible.

If a folder's configuration overrides the global settings then that configuration needs to be administered independently. This means that, when operational or capacity requirements change, in addition to updating the global settings, settings of all the independently administered folders must be updated as well. This increases the chance that the folder's configuration may be overlooked and end up configured in a way that is no longer appropriate. While there are certainly circumstances where overriding the global settings is the correct course of action, doing so does require additional time and attention on the part of the administrator.

8.2.2 Receiving Message Size

This controls whether this folder will use the default receiving message size limit (set in item 1.2.4) or whether the folder will observe different limits. If “Use default limit” is selected, the folder will observe the receiving message size limit set in the global policy in item 1.2.4. If “Maximum KB” is selected, the receiving message size limit as set in the global settings will be overridden and the text field on the right side of the panel will be used to hold the folder’s customized receiving size limit. See item 1.2.4 for recommendations regarding message size limits and the risks of various settings.

Note that the default Internet Newsgroups public folder and other folders that are not mail enabled do not have this tab or any of the controls contained therein. As such, configuring these folders is neither necessary nor possible.

It is recommended that, in general, folders use the setting prescribed in the global settings. There may be cases where it is reasonable or necessary to change the settings of a particular folder, but these cases should be limited as much as possible.

If a folder’s configuration overrides the global settings then that configuration needs to be administered independently. This means that, when operational or capacity requirements change, in addition to updating the global settings, settings of all the independently administered folders must be updated as well. This increases the chance that the folder’s configuration may be overlooked and end up configured in a way that is no longer appropriate. While there are certainly circumstances where overriding the global settings is the correct course of action, doing so does require additional time and attention on the part of the administrator.

8.2.3 Message Restrictions

This allows administrators to indicate that this folder should not receive mail from certain senders. It is used in addition to the limits imposed by items 1.2.2 (blacklist provider intended to block spammers) and 1.2.8 (intended to block certain senders on an enterprise wide basis). This means that, even if this control specifies that the folder is allowed to receive mail from “Everyone”, mail from senders given in items 1.2.2 or 1.2.8 will still be blocked. This field is simply intended to block mail from senders who should not be sending mail to this folder in particular. This control allows mail to be allowed from “Everyone”, from a limited set of users, or from everyone except a limited set of users. In addition, a further restriction may be applied to any of these settings in that only authenticated users may be allowed to send to this folder.

Note that the default Internet Newsgroups public folder and other folders that are not mail enabled do not have this tab or any of the controls contained therein. As such, configuring these folders is neither necessary nor possible.

In most cases this control can be set to “Everyone”, thus providing no additional limits on who may send mail to this public folder. In most cases where a sender should be blocked, the global level filters described in items 1.2.2 and 1.2.8 would be the correct place to specify the denied user. Some folders, however, will be intended for use only by internal (authenticated) users or by some other small set of users. In these cases, it would be reasonable to utilize this control to ensure that only those users with legitimate need to access this folder are given permissions to do so. As such, no recommendation is provided here.

If the list of users who are not permitted to send to this folder is too large then some users may not be able to perform legitimate activities involving this folder. Likewise, any restrictions on use of the folder will need to be reviewed periodically to make sure that users are not being unnecessarily denied (or granted) access to this folder. On the other hand, many folders should only permit a limited number of users to send mail to them. In these cases, the list should be tightly restricted to prevent unwanted mail to the folder.

8.3 Permissions

8.3.1 Client Permissions

This controls the rights that certain users and groups have to this public folder. The controls in this panel determine who is able to create, read, edit, delete, create sub-folders, and view the folder. The control panel contains a number of pre-defined roles with various commonly used combinations of access abilities, although any combination of permissions may be created by manually selecting the desired permissions.

Carefully scrutinize the permissions associated with each public folder and only include the most restrictive set of permissions needed to make use of the public folder in the intended manner. However, due to the wide variety of uses of public folders, no single recommendation for these settings can be given. Administrators must identify and grant required permissions on a folder-by-folder basis.

8.3.2 Directory Rights

This allows the administrator to control the mail-related actions on this public folder. These include the abilities to read, write, and delete the folder's contents. Users and groups may be assigned any of these rights.

The Directory Rights button will only be enabled if the directory is mail enabled. In general, it is recommended that folders not be mail enabled. If this recommendation is followed, the Directory Rights button is not enabled and no customization of these settings is necessary or possible. There are, however, cases where it could be reasonable to mail enable a folder. If this is done, carefully scrutinize the permissions associated with each public folder and only include the most restrictive set of permissions needed to make use of the public folder in the intended manner. Due to the wide variety of uses of public folders, no single recommendation for these settings can be given. Administrators must identify and grant required permissions on a folder-by-folder basis.

8.3.3 Administrative Rights

This allows the administrator to select which users will be able to use the Exchange System Manager or some other management interface to change various settings on this folder. Other settings governed by the control include modifying both the administrative (this item) and client ACLs (items 8.3.3 and 8.3.1, respectively) as well as storage limits (item 8.4.1), age limits (item 8.1.2), deletion settings (item 8.1.1), and other features. These capabilities can be assigned to both groups and individual users. Note that the "Special permissions" right will appear to be granted to most users and groups including the "Everyone" group. This right is required by Exchange, does not represent a security risk, and should not be changed.

Carefully scrutinize the permissions associated with each public folder and only include the most restrictive set of permissions needed to administer the folder. Due to the wide variety of uses of public folders, no single recommendation for these settings can be given. Administrators must identify and grant required permissions on a folder-by-folder basis.

8.4 Storage Limits

8.4.1 Storage Limits

This controls whether this public folder uses the default storage limits as set on the public folder store (items 7.7.1 and 7.7.2) or if the default should be overridden. The checkmark next to “Use public store defaults” controls whether the public store settings are overridden. If this checkbox is selected, then the public folder store’s settings are applied to this folder and the other fields in this group are disabled. If the checkbox is cleared, then the settings provided by the store are discarded and the remaining fields in this group (“Issue warning at”, “Prohibit post at”, and “Maximum item size”) are used to determine any storage limits for folder items. The latter fields function the same way as the corresponding fields (items 7.7.1 and 7.7.2) in the public folder store. Please consult these controls for recommendations and risks regarding this setting. Note that it is not possible to override just one of the three fields. If “Use public store defaults” is not selected then all the storage limit controls in the public folder store will be ignored for this folder. This means that, if the folder overrides the folder store and fails to enable one of the three features, then that feature will not be enabled on this folder regardless of the setting on the folder store. As such, be sure that the settings of all three fields reflect the desired actions if the public folder store’s configuration is overridden.

It is recommended that, in general, public folders use the setting prescribed by the public folder store. There may be cases where it is reasonable or necessary to change the settings of a particular folder, but these cases should be limited as much as possible.

If a folder overrides the settings of the public folder store then it needs to be administered independently. This means that, when operational or capacity requirements change, in addition to updating the settings on the store, settings of all the independently administered folders must be updated as well. This increases the chance that the folder may be overlooked and end up configured in a way that is no longer appropriate. While there are certainly circumstances where overriding the settings of the folder store is the correct course of action, doing so does require additional time and attention on the part of the administrator.

8.5 Visibility

8.5.1 Address List Name

This control allows you to specify the name of a given public folder as displayed to the user. The administrator can either specify a display name manually or can simply set the display name to the name of the folder as shown in the administrative console.

Note that the default Internet Newsgroups public folder and other folders that are not mail enabled do not have this tab or any of the controls contained therein. As such, configuring these folders is neither necessary nor possible.

The administrator should manually specify a name other than the name of the folder as given in the administrative console. This is called folder aliasing and is commonly used in web servers. The purpose of folder aliasing is to ensure that learning the name of the folder does not provide any information to a user as to how the folders are positioned on the server. By specifying a display name other than the folder's physical name, a user will not know how to refer to the folder outside the Exchange environment. This reduces the amount of information available to plan an attack on the public folder system.

It should be noted that, regardless of the setting of this control, the path of the folder will not be displayed to a user. This control is only used to hide the name of the folder itself. Using the name of the physical folder as the display name does not create any vulnerabilities in the system beyond giving users more information about how the system is configured than is strictly needed to use Exchange. However, it is relatively simple to perform and recommended that the folder aliasing feature be utilized.

8.5.2 Hide from Exchange Address List

This controls whether the email address of this public folder will appear on address lists published by Exchange. Note that this control does not govern who can send mail to this folder in any way – it only controls whether the address is made available via the Address Book by Exchange. Anyone who knows the address of the folder will be able to post to it regardless of this setting.

Note that the default Internet Newsgroups public folder and other folders that are not mail enabled do not have this tab or any of the controls contained therein. As such, configuring these folders is neither necessary nor possible.

It is recommended that this field be enabled (the checkbox be selected) to hide the folder's address. This reduces the distribution of this folder's address and reduces the opportunity for it to come into the hands of parties who might attempt to post spam or other irrelevant information to the folder. If the folder is intended to be publicly accessible then the setting of this control is less important since there would be other ways for interested parties to acquire its address. Most folders, however, will benefit from limited distribution of its address.

If this recommendation is not followed, the folder's administrators will need to be aware that the folder's address could end up being widely disseminated which may lead to spam or other irrelevant postings which will need to be deleted.

8.5.3 Send on Behalf

This controls whether other users can send mail on behalf of this public folder. For example, if the folder's email name is "A-folder" and the user "Bob" is added to this field in A-folder's properties, then Bob will be able to send mail on behalf of A-folder. The from field of the resulting message will look like:

From: Bob on behalf of A-folder

This can be useful if Bob is an administrator or moderator for A-folder and wishes to make it clear when he is sending messages in that capacity.

Note that the default Internet Newsgroups public folder and other folders that are not mail enabled do not have this tab or any of the controls contained therein. As such, configuring these folders is neither necessary nor possible.

It is recommended that the "Send on behalf" field be cleared (messages are not sent on behalf of any party). While the full from field displays both the actual sender as well as who the message is on behalf of, in many instances only the party on whose behalf the message was sent may be seen. For example, if the message is viewed only in the preview pane of Outlook, the actual sender may never be seen. As such, it is possible for a message's actual sender to be effectively hidden from the user using this feature.

If users are allowed to send on behalf of other parties, then users may never realize the actual sender of the message. This can allow senders to mask their activities. As such, if "Send on behalf" is used it is important that only trusted parties are given the ability to send on behalf of others, and their messages should be reviewed to ensure this privilege is not being abused.

9. Routing Groups

9.1 Authentication

9.1.1 Authentication and Encryption Algorithms

This panel controls the authentication and encryption algorithms used for outbound connections using this connector. (That is, the authentication used when delivering outbound mail to another SMTP Virtual Server.) It can be configured to Anonymous, Basic authentication (clear text), and Integrated Windows Authentication (also known as NTLM authentication). Only one authentication method may be selected. In addition to the authentication method, a checkbox allows TLS encryption to be utilized regardless of the authentication method selected.

Enable TLS encryption. Doing so prevents message content from being sent between SMTP servers in clear text. When the receiving end of the connector is locally controlled, Basic authentication should be used to authenticate the initiating end of the connection.

Failure to employ TLS encryption makes communication between SMTP Virtual Servers susceptible to eavesdropping. While sensitive message bodies should be encrypted by the sender at the client, using TLS provides added security for the body and also hides the message's senders and recipients from eavesdroppers. Failure to use TLS when Basic authentication is enabled is especially insecure since the authentication password would then be sent in the clear.

Failure to use authentication can allow mail server spoofing. While many attacks can be initiated regardless of whether SMTP servers authenticate each other, the authentication requirement does limit an attacker's ability to bypass filters and other defenses. Basic authentication is recommended over Integrated Windows Authentication (NTLM) because NTLM is a negotiated protocol and can result in variable levels of security. While NTLM servers can theoretically control this negotiation to only allow the most secure choices, this cannot be done through Exchange. (To only permit the most secure form of the protocol, NTLM v2, set the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\LMCompatibility registry key to 5. See Microsoft Knowledge Base Article 239869.)

Note that both authentication and encryption depend on the recipient of the outbound communication being capable of supporting these features. While this should be possible if both participants are controlled by the organization, communication with external SMTP virtual machines may not be securable in this way.

9.2 Connections

9.2.1 Connector Scope

This field controls which servers are permitted to use this SMTP Connector. The primary purpose of this control is to help limit use of this connector to computers for which it is the most efficient link. For example, if one group of computers had a fast connection to this connector and a second group had fast connections to a second connector, but the two groups were linked by a slow link, you might want to create separate routing groups for each of these groups and then enforce that each connector should only service the group to which it had the best connection. However, any control that limits allowed connectivity has obvious security ramifications as well.

Selecting “Entire Organization” allows any computer in the Exchange organization to use this connector. Selecting “Routing Group” means only those members of the connector's routing group may use the connector. Use of the connector should be limited to the Routing Group in order to limit and control general network connectivity.

Selecting “Entire Organization” might be necessary if it is impractical to create routing groups to efficiently subdivide the organization's network. However, doing this means that any party that is capable of appearing as a member of the organization will be able to use this connector. It also could complicate network monitoring due to the open connectivity of the connector.

9.2.2 Routing Options

This controls whether the SMTP connector routes its messages through simple DNS lookups (forming outbound connections to whichever host its routing table prefers) or whether it should use a Smart host. A Smart host is like an outbound proxy server and can serve as the clearing house for all outbound SMTP traffic. Note that setting this control to "Smart host" will override any settings placed in the virtual server's Smart host specification (item 11.2.4). For additional information on Smart hosts, please refer to the Exchange 2003 Admin Guide http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/admingde.mspx

Selecting "Smart host" means that all outbound messages through this connector will pass through a single server. This allows hardening to be applied to this single point rather than at multiple locations throughout your network. As such, the Smart host provides many of the advantages of a proxy server (and, indeed, could be the same machine).

The only exception to this recommendation is when one is configuring the external SMTP server. Since such a host would have no more outward facing computer to forward its messages to, it cannot be configured to use a Smart host, it should simply use DNS. Smart hosts may be chained together, but care must be taken that the chain is not cyclic or outbound mail will never travel beyond the Smart hosts.

If it is not practical to set up a Smart host then all potential outbound servers will need to be secured individually. This may prove challenging since the server's routing table may change due to external input, resulting in an outward path that was not previously expected. As such, potential outbound paths would need to be monitored regularly.

There are some potential issues if smart hosts are used at the virtual server level:

• If there is more than one computer running Exchange in your organization, mail flow between servers may not work.
• If an IP address is listed in the Smart host text box and is not enclosed in square brackets, mail flow may not work.
• If a name is listed in the Smart host text box, it must be a Fully Qualified Domain Name (FQDN) or mail flow between servers may not work.
For additional information, refer to the Microsoft Document entitled: "Exchange Server 2003 Transport and Routing" located at
http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/extransrout.mspx

9.3 Delivery Restrictions

9.3.1 Accept or Reject Outbound Messages through this Routing Group Connector

This controls whether the routing group connector accepts or rejects outbound messages from specific internal users or computers. It does not affect the delivery of inbound messages. Only users and computers known to Active Directory may be specified in the lists on this panel (but this is reasonable since one needs to be known in Active Directory to have an email account in Exchange). Note that only one exception list “Accept messages from/Reject messages from” is used at a time. Specifically, if the default action is “Accept”, only the “Reject messages from” list is used. If the default action is “Reject” only the “Accept messages from” list is used. Note that each individual routing group connector has its own, independent accept and reject lists. This means that if a particular user or computer is denied the use of this connector, they still may be able to use other routing connectors to send messages. If this is the case, the user will probably not even notice that a restriction was placed on their actions.

It is more secure to reject all messages by default and only permit messages to be sent through any given connector by explicitly specified users. However, even in the most secure environments, it seems likely that denying a user the ability to use a connector will be the exception rather than the rule. It is more likely that individual users or groups will be denied the ability to use a connector only in exceptional circumstances, and that specifying these users will not only be easier for the administrator to configure, but will also result in a policy that is easier to clearly understand using this control panel (since one will not need to run through a potentially long list of allowed users to determine whether or not a particular user has been denied access to this connector). Because of this no recommendation is made here.

Note that this only controls access to the routing group connector. If a user is denied the ability to use this routing group connector, they may (and probably should) have the ability to use some other routing group connector for their messages. This control is intended to allow for the creation of special purpose routing connectors which may only be used by a certain list of individuals, possibly to ensure that these individuals do not need to compete for network resources with other users, or for traffic compartmentalization purposes. To control a user's ability to send outbound messages globally, use sender filters (item 1.2.8).

9.3.2 Servers Allowed to Send Mail over Specific SMTP Connector

This is used to control which servers in this routing group will act as bridgehead servers. Bridgehead servers are used to handle connections between routing groups. Mail to one machine in the routing group will be sent to one of that group's bridgehead servers. The bridgehead server will then be responsible for passing the mail on to the correct server (possibly itself).

By default, all servers in a routing group act as bridgehead servers. This means that all servers in the routing group are capable of communicating with all external bridgehead servers directly. The administrator should specify the list of servers that are allowed to serve as bridgeheads for the routing group. This is done to better control connectivity between machines within the Exchange organization. This, in turn, makes it easier to monitor and debug network traffic. On the other hand, each routing group should have at least two bridgehead servers for increased availability (assuming the routing group consists of more than a single server). If only one bridgehead server is used and it fails, all email will fail and load balancing issues will increase. Select “These servers can send mail over this connector” even if every member in the routing group will be serving as a bridgehead anyway. This is done so that, if at some future point, another machine is added to this routing group, it will not immediately be promoted to acting as a bridgehead server without the administrator explicitly assigning it to this role.

Failure to implement this recommendation will allow open connectivity between the servers on this routing group and the bridgeheads of other routing group. This can result in a more complicated network topology which would make monitoring more difficult. It will also mean that a server that is added to this routing group will instantly become a bridgehead server, which may not be anticipated by the administrator.

9.3.3 Allow Messages to be Relayed over this SMTP connector to These Domains

This checkbox controls whether or not unauthenticated computers are allowed to resend (relay) email messages through this connector. (Authenticated users and computers can relay messages regardless of this control's setting.)

Selecting this checkbox on an Internet facing host would allow any party to use your Exchange Server to resend mail. This practice is often employed by spammers to obfuscate the source of their messages and can result in increased load on the mail server. This checkbox should remain cleared.

The only reason this checkbox might need to be selected is if you will need machines to relay through this server but cannot (or do not want to) configure them to authenticate themselves. If this is the case, do not use this connector to support a direct connection to the Internet. Otherwise, allowing unauthenticated relaying will almost inevitably result in abuse of the relay by spammers and increased load on the connector.

9.3.4 Accept or Reject Outbound Messages through this SMTP Connector

This controls whether the SMTP connector accepts or rejects outbound messages from specific internal users or computers. It does not affect the delivery of inbound messages. Only users and computers known to Active Directory may be specified in the lists on this panel (but this is not too surprising since one needs to be known in Active Directory to have an email account in Exchange anyway). Note that only one exception list “Accept messages from/Reject messages from” is used at a time. Specifically, if the default action is “Accept”, only the “Reject messages from” list is used. If the default action is “Reject” only the “Accept messages from” list is used. Note that each individual SMTP connector has its own, independent accept and reject lists. This means that if a particular user or computer is denied the use of this connector, they still may be able to use other SMTP connectors to send messages. If this is the case, the user will probably not even notice that a restriction was placed on their actions.

It would be more secure to reject all messages by default and only permit messages to be sent through any given connector by explicitly specified users. However, even in the most secure environments, it seems likely that denying a user the ability to use a connector will be the exception rather than the rule. It is more likely that individual users or groups will be denied the ability to use a connector only in exceptional circumstances, and that specifying these users will not only be easier for the administrator to configure, but will also result in a delivery policy that is easier to clearly understand using this control panel (since one will not need to run through a potentially long list of allowed users to determine whether or not a particular user has been denied access to this connector). Because of this, no recommendation is made here.

Note that this only controls access to the SMTP connector. If a user is denied the ability to use this SMTP connector, they may (and probably should) have the ability to use some other SMTP connector for their messages. This control is intended to allow for the creation of special purpose SMTP connectors which may only be used by a certain list of individuals, possibly to ensure that these individuals do not need to compete for network resources with other users, or for traffic compartmentalization purposes. To control a user's ability to send outbound messages globally, use sender filters (item 1.2.8).

9.3.5 Allowed Sizes for Routing Group Connector Messages

This allows the administrator to control the maximum size of outgoing messages on this connector. It is intended to help with flow control by ensuring that messages that exceed a certain size limit do not use certain connectors. For example, if one had a connector that represented a slow, expensive, or otherwise restricted connection, one might choose to restrict the size of messages that crossed the connector to ensure that large messages did not needlessly tie up the link. If the checkbox is cleared then no limits are applied and messages of any size may cross this link. If the checkbox is selected, the maximum size of allowed messages is set in the field on the right side of the panel. Note that the size limits applied to connectors are complimentary to size limits imposed by other items (1.2.4, 5.4.3, 7.7.2, 12.2.3, and so on). The settings of this control apply only to the associated connector. A large message will still be able to reach its destination as long as some connection exists that its size does not prohibit it from using. Only if no connector exists that permits messages of the given size to cross will the message fail to be delivered.

It is recommended that, in general, no limits are applied to connection sizes. This is done so that connectors do not end up prohibiting the delivery of messages that would otherwise be permitted by the Exchange configuration. If one wishes to prohibit messages of a certain size in the enterprise, the control appropriate to that message (items 1.2.4, 5.4.3, and so on) should be used. Using connectors to control size limits at an enterprise-wide level is discouraged since the limits must be applied to every potential connector in order to create an effective enterprise-wide limit. Such a scenario would be time consuming to set up and extremely difficult to debug if an error is made. There may be some situations where it is useful to limit message sizes for specific connectors for efficiency reasons, but if this is done care must be taken that some path still exists between the connected routing groups that does not limit message sizes.

If limits are applied to a connector, care must be taken that all messages will still be able to cross between the routing groups. This not only involves ensuring that larger messages can traverse between the routing groups, but also that messages of the appropriate type (item 9.3.8) and priorities can also connect between nodes. Failure to do this may result in an information flow failure that is extremely difficult to debug. As such, the utmost care should be taken when limiting message sizes on connectors.

9.3.6 Allowed Types for SMTP Connector Messages

This controls the types of communication permitted over this SMTP connector. The types that can be controlled here are “System messages”, which include messages originating from either Exchange or the Windows operating system (for example, replication messages, monitoring messages, and delivery reports), and “Non-system messages”, which include messages originating from users, groups, and contacts. It can be used to allow the administrator to create single purpose connectors that will only handle one type of traffic. This can be useful in that it can ensure that one type of message will not monopolize a particular connector at the expense of another message type. For example, if there are separate connectors for each type, then system messages will always be able use their connector even if the second connector becomes clogged with non-system messages.

Ideally, if SMTP connectors are used, one’s enterprise will contain multiple connectors between each SMTP server. If this is done, one or more connectors should be designated for the exclusive use of system messages. (One or more connectors should also allow non-system messages, although they need not be exclusive.) This ensures that system messages will have a route between SMTP servers even if a high volume of non-system messages monopolizes other connectors. Ultimately, however,, it is important to ensure that both system and non-system messages are able to travel between all SMTP servers.

If there is not a path between SMTP servers for a particular type of message then the functionality of the Exchange enterprise may be compromised. If both types of messages can be sent between SMTP servers, but the system messages must always share the connector with non-system messages, a high volume of non-systems messages may cause system messages to be delayed, possible resulting in performance problems.

9.3.7 Allowed Sizes for SMTP Connector Messages

This allows the administrator to control the maximum size of outgoing messages on this connector. It is intended to help with flow control by ensuring that messages that exceed a certain size limit do not use certain connectors. For example, if one had a connector that represented a slow, expensive, or otherwise restricted connection, one might choose to restrict the size of messages that crossed the connector to ensure that large messages did not needlessly tie up the link. If the checkbox is cleared then no limits are applied and messages of any size may cross this link. If the checkbox is selected, the maximum size of allowed messages is set in the field on the right side of the panel. Note that the size limits applied to connectors are complimentary to size limits imposed by other items (1.2.4, 5.4.3, 7.7.2, 12.2.3, and so on). The settings of this control apply only to the associated connector. A large message will still be able to reach its destination as long as some connection exists that its size does not prohibit it from using. Only if no connector exists that permits messages of the given size to cross will the message fail to be delivered.

It is recommended that, in general, no limits are applied to connection sizes. This is done so that connectors do not end up prohibiting the delivery of messages that would otherwise be permitted by the Exchange configuration. If one wishes to prohibit messages of a certain size in the enterprise, the control appropriate to that message (items 1.2.4, 5.4.3, and so on) should be used. Using connectors to control size limits at an enterprise-wide level is discouraged since the limits must be applied to every potential connector in order to create an effective enterprise-wide limit. Such a scenario would be time consuming to set up and extremely difficult to debug if an error is made. There may be some situations where it is useful to limit message sizes for specific connectors for efficiency reasons, but if this is done care must be taken that some path still exists between the connected SMTP servers that does not limit message sizes.

If limits are applied to a connector, care must be taken that all messages will still be able to cross between the SMTP servers. This not only involves ensuring that larger messages can traverse between the SMTP servers, but also that messages of the appropriate type (item 9.3.8) and priorities can also connect between nodes. Failure to do this may result in an information flow failure that is extremely difficult to debug. As such, the utmost care should be taken when limiting message sizes on connectors.

9.3.8 Allowed Types for Routing Group Connector Messages

This controls the types of communication permitted over this routing group connector. The types that can be controlled here are “System messages”, which include messages originating from either Exchange or the Windows operating system (for example, replication messages, monitoring messages, and delivery reports), and “Non-system messages”, which include messages originating from users, groups, and contacts. It can be used to allow the administrator to create single purpose connectors that will only handle one type of traffic. This can be useful in that it can ensure that one type of message will not monopolize a particular connector at the expense of another message type. For example, if there are separate connectors for each type, then system messages will always be able use their connector even if the second connector becomes clogged with non-system messages.

Ideally, one’s enterprise will contain multiple connectors between each routing group. If this is done, one or more connectors should be designated for the exclusive use of system messages. (One or more connectors should also allow non-system messages, although they need not be exclusive.) This ensures that system messages will have a route between routing groups even if a high volume of non-system messages monopolizes other connectors. Ultimately, however, it is important to ensure that both system and non-system messages are able to travel between all routing groups.

If there is not a path between routing groups for a particular type of message then the functionality of the Exchange enterprise may be compromised. If both types of messages can be sent between routing groups, but the system messages must always share the connector with non-system messages, a high volume of non-systems messages may cause system messages to be delayed, possible resulting in performance problems.

10. Servers

10.1 Deletion Settings

10.1.1 Zero Out Deleted Database Pages

This feature controls how deleted memory is handled. If this feature is not enabled, when a mail message or public folder posting is deleted (this is to say, all references to it are removed and the management program actually deletes the message) the operating system simply marks the memory that was previously used to store the message as available for use. Eventually, this memory may be utilized to store additional information. However, until this time, the message will still be present on the disk and certain utilities will be able to recover some or all of the message. Since legitimate administrators should have access to backup records of the message, this method of message retrieval would be primarily of use by attackers to reconstruct deleted data. If the "Zero out deleted database pages" checkbox is selected, before the disk memory is released back to the operating system, it is overwritten with null data. This means that by the time the memory is returned to the operating system, it essentially no longer contains any information that would allow the message to be retrieved.

Select the "Zero out deleted database pages" checkbox to enable this feature. This will limit the chance of unauthorized users being able to retrieve deleted material. The only disadvantage to this feature is that message deletion will be more time consuming since the server must actually overwrite the entire message. However, since the administrative process that handles message deletion usually runs during off hours, it is unlikely that this slowdown will be visible to users.

If this feature is not enabled then extra care must be taken that access to the server's hard drive (both network and physical access) is limited to trusted individuals. This includes limiting regular, unprivileged logon access. Aside from a slight slowdown during nightly message deletion, there is no reason not to enable this feature. Administrators should not view disabling this feature as providing a way to retrieve deleted information since attempting to pinpoint a specific message in memory will be very difficult, and it will be entirely a matter of chance as to whether the operating system has already utilized the message's memory and overwritten some or the entire message. It is much better to use traditional backup mechanisms for recovery of lost material and to zero out deleted messages on the server to keep the hard disk secure.

10.2 Display

10.2.1 Server Policy Display

Server Policies (created and stored under the Exchange System Manager → Administrative Groups → [administrative group] → System Policies container) can be used to control the settings configured in the server container's "General" Tab. A Server Policy can be applied to multiple servers in this administrative group, obviating the need to configure the servers individually and ensuring a uniform configuration across the various servers. If multiple servers are used, and if these servers are intended to have the same configuration (at least on the General Tab) then the use of Server Policies is recommended both to simplify administration and to ensure configuration consistency.

This screen simply lists the Server Policies that have been applied to this server. Use it to verify that a policy object has been successfully applied to the server. The policy itself cannot be manipulated through this interface.

10.3 Logging

10.3.1 Enable Circular Logging

This controls how log files are written. If circular logging is enabled, there is one log file for this storage group with a maximum size of 5MB. Once the 5MB limit has been reached, additional log entries begin overwriting the oldest log entries. The result is that the log file contains all the most recent audit events without exceeding 5MB in size. If circular logging is disabled, once a log file reaches 5MB a new log file is created. This log file is then used until it reaches 5MB at which point yet another log file is started, and so on. These log files remain on the server until manually deleted.

Clear the "Enable circular logging" checkbox. While circular logging is more memory efficient, it can result in log events getting overwritten before they are backed up or analyzed. Disabling circular logging, while potentially memory-expensive, can ensure that log entries are retained as long as necessary. Log files should be backed up and retained for at least a month (or longer depending on organizational policy). Once logs have been copied to backup or other offline medium, they can be deleted from the Exchange server, thus freeing up space.

The only valid reason to enable circular logging is if the Exchange server is so low on memory that even a small number of log files could potentially fill the disk and render the server inoperable. If this occurs, circular logging could be enabled as a temporary fix while additional hard disk space is added to the server. While circular logging is enabled, backups of the log file should be made frequently to prevent the possibility of log entries being permanently lost. Note that the use of circular logging on backend servers may prevent the backup service from being able to restore the server completely due to lost data when the log overwrites itself.

Note that, on front-end servers, the only messages sent to the log file are simple status messages for routine tasks. As such, they are of little interest to administrators. For this reason, on front-end servers, circular logging should be enabled to prevent the front-end server from filling with useless log information.

10.3.2 Diagnostic Log Levels

These control the collection of diagnostic information on the Exchange server. Logging is broken up into 14 main "services" each of which has anywhere from 2 to 26 "categories" of events to be monitored. Moreover, each category may be set to one of four levels of logging: None (logging disabled), Minimum, Medium, and Maximum, depending on how much detail one desires. The higher the level of detail, the more memory required to store the audit material.

For general operation, all categories of all services should be set to "None". Diagnostic logging is intended to help administrators debug problems with their systems, not as a general purpose auditing tool. The diagnostic logs collect a great deal of information – in one trial, setting a few categories to Maximum resulted in a GB of data in only a few hours. As a result, diagnostic log files can grow huge very quickly and, from a security perspective, have a very poor signal-to-noise ratio. (That is, for every piece of security relevant information, an administrator might need to wade through hundreds or thousands of unrelated entries.) Given that other auditing facilities exist that are better targeted to security events (items 2.5.1, 11.4.1, and so on), there is no reason to use diagnostic logs on a regular basis. Instead, specific categories can be enabled for limited periods of time when attempting to debug relevant pieces of Exchange functionality. Once debugging has finished, diagnostic logging should be disabled again.

If diagnostic logging is enabled, administrators will need to monitor log file growth and remove unnecessary log files quickly. Failure to do so can rapidly deplete system memory.

10.3.3 Enable Subject Logging and Display

When message tracking is enabled (item 10.3.4), by default only the sender, recipients, time, and other delivery information is included. Information that does not relate directly to delivery functions, such as the subject and message body, is not included. However, the absence of the message subject line can make it difficult to locate a specific message in the log unless one knows roughly what time the message was sent. To simplify searches through these logs, Exchange offers the ability to include the message subject line in the log files and in the Message Tracking Center display. This can make it significantly easier to locate a specific message at the cost of larger log files.

The disadvantage of this feature is that log files increase in size to accommodate the subject information. In addition, the storage and browsing of subject information may raise privacy and legal concerns - enterprise policy should be consulted before this feature is enabled. Finally, since the log files may contain sensitive information in the form of the subject line, the log files will need to be protected, both on and off the server. (Protection of the log files that only include delivery information is less of an issue since, without the mail archive, the subjectless logs would have less value to an attacker.) For these reasons, it is recommended that subject logging not be enabled. The tradeoff of this is that finding the correct message in the message tracking logs will become more difficult since the administrator will need to search using only the time the message was sent and the message's sender.

The direct security risks of enabling this feature (namely the possibility that sensitive information in the subject line could be disclosed to unauthorized individuals) are fairly slight. The larger issue is that logging subjects can significantly increase the size of each entry in the message tracking log. As most enterprises likely process a great deal of mail to begin with, these logs are already likely to grow rapidly and adding subject information exacerbates this problem significantly. As a result, administrators who enable subject logging need to be especially careful about the size of the log files to prevent the logs from filling the partition.

Note that if the recommendation in 10.3.4 is followed (message tracking is not enabled) then the setting of this control will have no practical effect on the system. That said, subject logging should still be disabled in case message tracking is turned on at a later time.

It is required that all log files be written to separate partitions from the partitions used by the Exchange Stores. Exchange will dismount its stores if it detects that it has run out of disk space, resulting in a complete loss of Exchange services. To minimize the chance of this happening, log files should write to some other partition so that if the logs fill this partition it will not result in the failure of Exchange.

10.3.4 Enable Message Tracking

The Exchange server has the ability to record delivery information for all messages that pass through this Information Store (in either direction). By enabling message tracking, a log is created that records what happens to a message while it is in the Exchange network (that is, from the time it arrives on the server to the time that it is handed off to an external server). A new log file is created each day. These log files do not contain the message body - just routing information. This log can be viewed and searched using the Exchange System Manager → Tools → Message Tracking Center.

On very active mail servers, these log files may grow fairly large, fairly quickly. Because of this, in general message tracking should be disabled. It should only be turned on by a knowledgeable administrator who knows exactly how to handle the enormous volume of data. Additional information can be found at:
http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/ex2k3ad.mspx

If message tracking is enabled, administrators should monitor its size to ensure that log files do not monopolize disk space. It is required that all log files be written to separate partitions from the partitions used by the Exchange Stores. Exchange will dismount its stores if it detects that it has run out of disk space, resulting in a complete loss of Exchange services. To minimize the chance of this happening, log files should write to some other partition so that if the logs fill this partition it will not result in the failure of Exchange.

10.3.5 Remove Log Files (days)

This feature allows the administrator to automate the task of deleting log files created by the message tracking feature (item 10.3.4). As noted, this feature creates a new log file every day. By enabling this feature and specifying the maximum age of log files in days, all log files that are older than this maximum will be deleted.

Clear the “Remove log files” checkbox. All log files should be retained for at least a month, and on a busy server, a month's worth of message tracking logs, especially if the logs include message subjects (item 10.3.3), will create a sizable burden on the server's memory usage. It is advisable instead to copy older log files to an external storage area and then delete the copied logs from the server. Note that if the recommendation in item 10.3.4 is followed (message tracking is not enabled) then the setting of this control will have no practical effect on the system. That said, automatic deletion should still be disabled in case message tracking is turned on at a later time.

The only disadvantage to disabling this feature is that, if for some reason, the logs are not regularly archived or deleted, the memory used to store the logs on the server may grow prohibitively. However, the best remedy for this remains backing up the logs and deleting them rather than having the server blindly deleting log files and risking loss of important information.

10.4 Monitoring and Tools

10.4.1 Monitoring Information

This screen displays the status of the SMTP queues on the Exchange server. There is nothing security related that can be set in this window, but it provides valuable information for system administrators and should be reviewed periodically to make sure all the queues are operating correctly. Of the various fields displayed, the State field is of most interest: if any queues have values other than Ready or Active, then there is probably an error with the system that needs to be addressed quickly to ensure service availability. This panel also contains the "Disable Outbound Mail" button that allows the administrator to instantly disable all outbound messages should that become necessary.

There are several queues listed in this screen. They include (but may not be limited to):

- Local delivery - Holds inbound messages to be delivered to a mailbox on this server.
- Messages awaiting directory lookup - Holds inbound messages whose recipients have not yet been identified in Active Directory.
- Messages pending submission - Holds messages that have been accepted by the SMTP Virtual Server but which have not yet been processed.
- Messages waiting to be routed (2 queues - SMTP & X400) - Holds outbound messages awaiting routing resolution of the next machine to which they should be sent.
- DSN messages pending submission - Holds Delivery Status Notification messages (non-delivery reports, and so on) awaiting processing.
- Failed message retry queue - Holds outbound messages whose previous attempt(s) at delivery were unsuccessful but for which additional attempts will be made. (If this queue grows suddenly, it may represent a problem with Internet connectivity.)

In addition, there will be a queue for each mailbox store (holding messages going into and out of the store) as well as ephemeral queues for each domain (that is mitre.org, yahoo.com, and so on) to which outbound mail is currently being sent.

Monitoring the mail queues can be used to help tune various routing controls in Exchange, as well as to help the administrator pin down problems with mail delivery.

10.4.2 Display Domain Controllers

This screen displays the list of known Domain Controllers (Configuration, Standard, and Global Catalog Servers) in the Exchange organization. Administrators may use this panel to verify that all appropriate Domain Controllers can be found by the current server. This can be used to diagnose connectivity issues between servers in the organization.

This screen cannot be used to control servers or their connectivity - it is simply present to display the current set of known Domain Controllers.

At the bottom of the panel is a checkbox that allows the administrator to control whether Domain Controllers are automatically discovered or manually added to the list of controllers. This checkbox should remain selected (enable automatic discover of servers) since failure to do this defeats the primary advantage of this panel, which is to inform the administrator of which Domain Controllers the server can detect.

10.4.3 Default Microsoft Exchange Services

The details button of this panel can be used to view the current status of the various monitored resources. By default only “Default Microsoft Exchange Services” are monitored, but other resources may be monitored in this panel as well items 10.4.6 through 10.4.12. This information can be used by administrators to monitor server behavior.

The various detail screens can also be used to configure when warning and critical alerts are sent out as described in items 10.4.6 - 10.4.13. However, no other administrative activities are possible in this panel.

The same panel can be reached using the path: Exchange System Manager → Tools → Monitoring and Status → Status → Properties [server] → Monitoring Tab → Detail button

10.4.4 Automatically Send Fatal Service Error Information to Microsoft

This feature controls whether debugging messages are sent to Microsoft whenever system errors are detected. Exchange automatically detects many system errors and can report them and additional debugging data to Microsoft, who can then use this information to improve the robustness of their product.

Disable this feature in your Exchange organization. While debugging information should never contain sensitive information, at the very least it can alert eavesdroppers to the existence of problems in your Exchange organization. This will, at the very least, alert them as to possibly advantageous timing to mount an attack. At worst, it may even provide them with information as to which aspects of Exchange are causing problems and might be vulnerable (or at least sensitive) to attack.

Failure to follow this recommendation means that all system errors in Exchange will result in external traffic that may be identified by an eavesdropper. This is, admittedly, a fairly minor risk and the feature is helpful to Microsoft in the production of more reliable products so some administrators may still wish to enabled it. For more secure environments, however, disabling this feature is recommended.

10.4.5 Disable all Monitoring of This Server

This control allows the administrator to disable all monitoring processes on this server. This reduces the amount of processing required by the server. The disadvantage of selecting this option is that it disables virtually all of Exchange's built in safety checks to warn the administrator in the case of malfunctions.

Monitoring should never be disabled on the server. The added processing cycles required to run the monitoring features should not be noticeable on an adequately-sized machine. Additionally, the information provided by the monitoring tools provides the first line of warning in the case of Exchange component failure, thereby increasing availability.

10.4.6 Warning State/Critical State of Disk Space Threshold

This field allows the administrator to control when each type of notification trigger is issued in response to low disk availability. By default, this resource is not added - it must be added manually before notification triggers can be enabled. The type of notification trigger is then used to determine the type of automated notification that occurs. Setting this control has no effect unless a notification is set up (using item 10.4.13). There are two types of notification triggers: "Warning" and "Critical".

This resource monitor first needs to be created by clicking the Add button and selecting "Disk Space Threshold". One monitor should be added for each disk used by the server. Once the monitor has been created, set the warning state to 3000 (MB remaining), and the critical state to 2000 (MB remaining). There should always be around 3 GB of available space on any of the disks that Exchange writes to. If there is regularly less space than this, more and/or larger disks may be needed. If the disk space falls below 2 GB, Exchange is getting close to running out of space. Immediate administrative action (backing up and removing log files, and so on) is needed to ensure that the server remains functional. On larger disks, or Exchange servers that expect faster storage utilization, these numbers might need to be increased to provide more time to respond to storage issues. While the warning triggers should never be less than the values given above, a good rule of thumb is to issue warnings when free space falls under 15% and critical messages when it falls under 5%.

Monitor disk availability. If the server were ever to run out of disk space, the server could fail catastrophically, possibly with data loss. The precise levels at which warning and critical notifications are dispatched can vary, but regular warning notifications should alert an administrator to the need to mitigate the load on the server, while critical notifications should prompt immediate action. For more info please refer to http://www.microsoft.com/exchange/downloads/2003/exbpa/default.asp

10.4.7 Duration / Warning State / Critical State of CPU Utilization

This field allows the administrator to control when each type of notification trigger is issued in response to high CPU utilization. By default, this resource is not added - it must be added manually before notification triggers can be enabled. The type of notification trigger is then used to determine the type of automated notification that occurs. Setting this control has no effect unless a notification is set up (using item 10.4.13). There are two types of notification triggers: "Warning" and "Critical".

This resource monitor first needs to be created by clicking the Add button and selecting "CPU Utilization Threshold". Once the monitor has been created, set the duration to 10 (minutes), the warning state to 80 (percent used), and the critical state to 90 (percent used). CPU utilization should only exceed 80 occasionally. If it starts doing so on a regular basis, then the server is either running too many processes and/or needs to be upgraded. Should the CPU utilization ever exceed 90 percent for more than a few minutes, the server is too heavily loaded and is in danger of being unable to service user requests.

The recommendation outlined here assumes that the average CPU load on a server is between 50 and 60 percent, with peaks approaching 80. If this is not the case, the warning and critical thresholds may need to be adjusted. Monitor CPU utilization since the CPU utilization plays a major part in the speed with which requests are processed. The precise levels at which warning and critical notifications are dispatched can vary, but regular warning notifications should alert an administrator to the need to mitigate the load on the server, while critical notifications should prompt immediate action.

10.4.8 Warning State / Critical State of SMTP Queues

This field allows the administrator to control when each type of notification trigger is issued in response to prolonged growth of SMTP queues. By default, this resource is not added - it must be added manually before notification triggers can be enabled. The type of notification trigger is then used to determine the type of automated notification that occurs. Setting this control has no effect unless a notification is set up (using item 10.4.13). There are two types of notification triggers: “Warning” and “Critical”.

This resource monitor first needs to be created by clicking the Add button and selecting “SMTP Queue Threshold”. Once the monitor has been created, set the warning state to 10 (minutes of continuous growth), and the critical state to 20 (minutes of continuous growth). Continuous queue growth for more than 10 minutes should happen only occasionally. If it happens more often, the server may have bandwidth problems and may need to be upgraded. Continuous growth for more than 20 minutes is generally indicative of a network problem and will need to be addressed immediately.

Monitor SMTP queue growth since anomalies in this resource directly impact mail delivery. The precise levels at which warning and critical notifications are dispatched can vary, but regular warning notifications should alert an administrator to the need to mitigate the load on the server, while critical notifications should prompt immediate action. Normal network activity (such as sending a large number of non-delivery reports due to a spammer hitting your server, or due to the presence of periodic connectors which only send mail at specific times and would queue messages at all other times) may result in notifications being sent. Administrators should take these into consideration when setting thresholds and responding to events.

10.4.9 Warning State / Critical State of X.400 Queue Thresholds

This field allows the administrator to control when each type of notification trigger is issued in response to prolonged growth of X.400 queues. By default, this resource is not added - it must be added manually before notification triggers can be enabled. The type of notification trigger is then used to determine the type of automated notification that occurs. Setting this control has no effect unless a notification is set up (using item 10.4.13). There are two types of notification triggers: “Warning” and “Critical”.

This resource monitor first needs to be created by clicking the Add button and selecting “X.400 Queue Threshold”. Once the monitor has been created, set the warning state to 10 (minutes of continuous growth), and the critical state to 20 (minutes of continuous growth). Continuous queue growth for more than 10 minutes should happen only occasionally. If it happens more often, the server may have bandwidth problems and may need to be upgraded. Continuous growth for more than 20 minutes is generally indicative of a network problem and will need to be addressed immediately.

Monitor X.400 queue growth since anomalies in this resource directly impact mail delivery. The precise levels at which warning and critical notifications are dispatched can vary, but regular warning notifications should alert an administrator to the need to mitigate the load on the server, while critical notifications should prompt immediate action. Normal network activity (such as sending a large number of non-delivery reports due to a spammer hitting your server, or due to the presence of periodic connectors which only send mail at specific times and would queue messages at all other times) may result in notifications being sent. Administrators should take these into consideration when setting thresholds and responding to events.

10.4.10 Duration / Warning State / Critical State of Virtual Memory Threshold

This field allows the administrator to control when each type of notification trigger is issued in response to low virtual memory. By default, this resource is not added - it must be added manually before notification triggers can be enabled. The type of notification trigger is then used to determine the type of automated notification that occurs. Setting this control has no effect unless a notification is set up (using item 10.4.13). There are two types of notification triggers: "Warning" and "Critical".

This resource monitor first needs to be created by clicking the Add button and selecting "Virtual Memory Threshold". Once the monitor has been created, set the duration to 3 (minutes), the warning state to 25 (percent remaining), and the critical state to 10 (percent remaining). Virtual memory should generally not fall below 25 for more than a couple of minutes. If it starts doing so on a regular basis, then the server is running too many processes and/or needs to be upgraded. Should virtual memory drop below 10 percent for more than a few minutes the server is too heavily loaded and is in danger of being unable to service user requests.

Monitor virtual memory since the virtual memory plays a major part in the speed with which requests are processed. The precise levels at which warning and critical notifications are dispatched can vary, but regular warning notifications should alert an administrator to the need to mitigate the load on the server, while critical notifications should prompt immediate action. For more info please refer to http://www.microsoft.com/exchange/downloads/2003/exbpa/default.asp

10.4.11 When Windows 2000 Service is Not Running, Change State to

This field allows the administrator to control the type of notification trigger issued when a selected Windows 2000 service stops. By default, this resource is not added - it must be added manually before notification triggers can be enabled. The type of notification trigger is then used to determine the type of automated notification that occurs. Setting this control has no effect unless a notification is set up (using item 10.4.13). There are two types of notification triggers: "Warning" and "Critical".

This resource monitor first needs to be created by clicking the Add button and selecting "Windows 2000 service". When creating the new monitor, add the following services: Critical; add Event Log, NTLM Security Support Provider, Remote Procedure Call (RPC), Server, Workstation, IIS Admin Services, and HTTP SSL. If HTTP, IMAP, POP3, or NNTP are used on the server, add the HTTP SSL, Microsoft Exchange IMAP4, Microsoft Exchange POP3, or Network News Transfer Protocol (NNTP) services, respectively. The core Exchange services (monitored in item 10.4.12) depend on these services to function, so a failure in any of these will lead to failure of one or more Exchange services.

Once all the above services have been added, the "When service is not running change state to" field should be set to Critical. The trigger should be "Critical" because, if any of the services that the core Exchange services depend on stop, this will require immediate attention.

If this field is not set to Critical it will be set to Warning. If this is done, warning notifications will need to be handled immediately, as well as critical notifications.

10.4.12 When Core Exchange Services are Not Running, Change State to

This field allows the administrator to control the type of notification trigger issued when any of the core Exchange services stops. The type of notification trigger is then used to determine the type of automated notification that occurs. Setting this control has no effect unless a notification is set up (using item 10.4.13). There are two types of notification triggers: “Warning” and “Critical”.

This field should be set to Critical since, if any of the core Exchange services stop, this will require immediate attention.

If this field is not set to Critical it will be set to Warning. If this is done, warning notifications will need to be handled immediately, as well as critical notifications.

10.4.13 Notification Creation

This container stores notifications. Notifications can either send out messages or run scripts if a warning or critical trigger is dispatched by a resource monitor (items 10.4.6 - 10.4.12). Notifications are configured to respond to one of the two triggers.

Create at least one notification for each type of event trigger. Moreover, notifications responding to critical triggers should include some way to immediately notify an administrator (via scripts that result in pager or phone alerts, and so on) as a simple email message may not be noticed until too late. Configure notifications on servers other than the server that is being monitored (the “Servers and connections to monitor” field). The reason for this is that, if a server is handling notifications for its own monitors, and if the server suddenly enters a critical state, it may be unable to send out notifications. Servers should send out notifications for each other as they cannot reliably do so for themselves. A reasonable way to do this would be to configure each exchange server to monitor “All servers” for both types of triggers as this provides redundancy. It is possible to have multiple notifications issued for a single trigger (for example, two emails and one script may be run in response to a single critical trigger), however, ensure at least one notification should exist for each of the warning and critical triggers.

The notification creation panel also gives administrators the ability to monitor connectors as well as servers. Connector events are automatically configured and do not have critical or warning levels. Connector triggers are created when a connector ceases to function for some reason. Notifications should be created for all connectors as well.

Failure to create a notification for each trigger will mean that, even if a monitor detects that a resource has entered a state that should trigger a warning or critical message, no alert will be dispatched to administrators. This can result in a failure to detect a problem until too late.

11. SMTP

11.1 Authentication

11.1.1 Authentication Method Used for Access to SMTP Virtual Directory

This controls the form of authentication used when attempting to connect to this virtual server. Possible settings are “Anonymous Access”, “Basic authentication” and “Integrated Windows Authentication”. If Basic authentication is selected, an additional checkbox allows the server to require that SSL/TLS encryption be used to protect the password (which would otherwise be sent in clear text). Integrated Windows Authentication is the Microsoft NTLM protocol. Note that if users must authenticate to the server (that is, Anonymous Access is cleared) the “Users” button may be used to grant or deny access to the server to specific users, should the administrator wish to do this.

If this virtual server will only communicate internally, Basic authentication and Require SSL/TLS should be selected in this panel. NTLM can be secured, but it is a negotiated protocol and can lead to variable levels of security. The use of SSL/TLS not only protects the username and password during authentication, but encrypts the mail messages as they are being transmitted, preventing eavesdroppers from reading messages.

Failure to implement SSL/TLS when Basic authentication is enabled is extremely insecure as it results in the user's identity and password being sent over the network in clear text. If this happens, network sniffers can easily collect these credentials. The use of NTLM (Integrated Windows Authentication checkbox), while it can protect the username and password during authentication, it does not provide encryption of message bodies, potentially allowing message content to be sniffed over the wire. Moreover, NTLM negotiates the details of how it secures the authentication and some possible selections are insecure. While it is possible to configure NTLM server to only use the more secure options, Exchange does not provide this capability. (To only permit the most secure form of the protocol, NTLM v2, set the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\LMCompatibility registry key to 5. See Microsoft Knowledge Base Article 239869.)

If the virtual server is intended to receive messages from external SMTP Virtual Servers, Anonymous Access will also need to be selected since external servers will not have the necessary credentials to authenticate to the server.

11.1.2 Resolve Anonymous Email

This feature causes the server to use a Directory Naming Service (DNS) lookup to try to determine the source of anonymous email at the cost of a small amount of additional processor use per anonymous message.

Although enabling this feature does not pose a security hazard, in general it should be disabled as a waste of resources. Anonymous email is invariably spam and should be filtered by the server immediately (item 1.2.10). Moreover, the DNS lookup is not necessarily a reliable indicator of perpetrator information because the source of the messages could be a compromised email account or the DNS entries could have been maliciously altered. As such, any result that the lookup returns is unlikely to be of any use to the mail's recipient. Given this, the extra processing required by this feature does not provide any benefit and should be avoided.

11.1.3 Certificate Wizard

This button starts a wizard to install a certificate to be used by this SMTP Virtual Server. Server certificates are required for many security features in Exchange, and without them the server cannot engage in many forms of secure communication. The wizard can guide the administrator through the process of requesting a new certificate or of importing an existing certificate. Certificates must be manually installed on each virtual server. This means that installing a certificate on one SMTP Virtual Server does not give other SMTP Virtual Servers (or virtual servers of any other protocol) access to this certificate. However, once a certificate is installed on one virtual server, any other virtual server (regardless of protocol used) may easily be configured to use this certificate by selecting “Assign an existing certificate” in the first page of the Wizard.

Install certificates on this virtual server. Without it, many other recommendations in this document concerning secure communication will be impossible. For highest security assurance, each virtual server should have its own certificate that it does not share with other servers. This reduces the damage due to server compromises and provides per-server identification.

Failure to implement this recommendation makes it virtually impossible to secure Exchange's communications. Use of any virtual server that has not been given a certificate should be considered a highly insecure action.

11.1.4 Authentication Method Used for Outbound Connections

This controls the authentication and encryption used for outbound connections using this virtual server. (That is, the authentication used when delivering outbound mail to other SMTP Virtual Servers.) It can be configured to Anonymous, Basic authentication (clear text), and Integrated Windows Authentication (also known as NTLM authentication). Only one authentication method may be selected. In addition to the authentication method, a checkbox allows TLS encryption to be utilized regardless of the authentication method selected.

Enable TLS encryption. Doing so prevents message content from being sent between SMTP servers in clear text. When the receiving server is locally controlled, Basic authentication should be used to authenticate the initiating end of the connection.

Failure to employ TLS encryption makes communication between SMTP Virtual Servers susceptible to eavesdropping. While sensitive message bodies should be encrypted by the sender at the client, using TLS provides added security for the body while also hiding the message senders and recipients from eavesdropping.

Both authentication and encryption depend on the recipient of the outbound communication being capable of supporting these features. While this should be possible if both participants are controlled by the organization, communication with external SMTP virtual machines may not be securable in this way. As a result, since using non-Anonymous authentication would likely end up interrupting the delivery of mail to external sources, this control should be configured to use Anonymous authentication. Note also that this control governs all outbound communication from this virtual server uniformly. If the server is in a mixed environment (some recipient servers support authentication and encryption while others do not) then this control will not be capable of specifying different actions for different servers. Under such circumstances, use the outbound security of SMTP connectors to individual servers to configure outbound authentication and encryption on a per-server basis (item 9.1.1).

11.2 Connections

11.2.1 Require Secure Channel (SSL) and Require 128 bit encryption to SMTP Virtual Server

This controls whether client machines should be forced to use secure channels to communicate with this virtual server. If this feature is enabled, clients will only be able to upload mail to the server if they are capable of supporting secure communication with the server. If secure channels are required, the server can also require the channel to be strongly secure by requiring 128 bit encryption.

Require secure channels and 128-bit encryption. The use of secure communication prevents eavesdroppers from reading or modifying communications between servers and clients. While sensitive message bodies should be encrypted by the sender at the client, securing the communication between the server and the recipient provides an added level of security while also protecting sender and recipient information that cannot be encrypted by the sender.

If a particular client application in your organization does not support secure communication, this feature will need to be disabled or that client will be unable to connect to the server. However, it is strongly advised that any such clients be upgraded. Failure to enable this feature gives attackers more opportunities to read and modify communication between the client and server. Requiring secure channels but not requiring 128-bit encryption will still protect most messages from eavesdropping and modification, but makes it slightly more likely that a determined and well equipped attacker may be able to overcome the channel security.

11.2.2 Exclude or Limit Relaying SMTP Servers

SMTP relays are used to pass messages through an SMTP server to another destination on the Internet. Under many circumstances relays are not needed as the SMTP Virtual Server may be configured to put mail on the Internet without the use of an intermediary. However, if an SMTP Virtual Server does not have a direct connection to the Internet, such as if it is in a protected subnet, then it will need to use another machine that does have a direct connection to the Internet (such as a front-end server) as a relay.

SMTP relays must be protected. If adequate security is not provided for the relay functionality then third parties may be able to use your relay service for their own uses. Most commonly, hijacking of relays is done by spammers to disguise the source of their messages, although vulnerable relays may also be used to cover the source of more destructive attacks. This not only results in unnecessary load on a server that is acting as the relay, but it potentially opens the possibility of legal liability. As such, “Only list below” should be selected and the allowed SMTP servers should be added explicitly.

This control is used to limit the servers that may use this server as a relay. By default, the control is configured so that only specified computers are allowed to relay through this server (with the caveat that, by default, the control described in item 11.2.3 is enabled, which overrides the setting of this control). Since the list of specified computers is blank, no machines should be able to relay through this machine (again, with the caveat of item 11.2.3). The control can be set to allow all computers to relay through the server except for a given list of machines, but this should almost never be selected. Regardless of how many entries are listed in a deny list, there will always be additional malicious computers that will have been missed. In virtually all cases, only a finite (possibly empty) list of selected computers should be allowed to use this server as a relay.

11.2.3 Exclude or Limit Relaying SMTP Servers

SMTP relays are used to pass messages through an SMTP server to another destination on the Internet. Under many circumstances relays are not needed as the SMTP Virtual Server may be configured to put mail on the Internet without the use of an intermediary. However, if an SMTP Virtual Server does not have a direct connection to the Internet, such as if it is in a protected subnet, then it will need to use another machine that does have a direct connection to the Internet (such as a front-end server) as a relay.

SMTP relays must be protected. If adequate security is not provided for the relay functionality then third parties may be able to use your relay service for their own uses. Most commonly, hijacking of relays is done by spammers to disguise the source of their messages, although vulnerable relays may also be used to cover the source of more destructive attacks. This not only results in an unnecessary load on a server that is acting as the relay, but it could potentially increase the possibility of legal liability.

This control is used to limit the servers that may use this server as a relay. If selected, any computer that can successfully authenticate to this server's domain may relay through it, potentially overriding settings provided in the list at the top of the panel (item 11.2.2). It is recommended that “Allow all computers which successfully authenticate, regardless of the list above” be selected. Doing this should prevent unknown servers from relaying through this server while ensuring that known servers retain access to it. If certain servers should never legitimately serve as relays then the administrator may wish to clear this control while leaving the list of allowed relaying servers in item 11.2.2 blank. Doing this will effectively prohibit relaying through this server.

If this control is cleared (authentication does not automatically allow relaying through this server) and this server should be used as a relay, the servers that should relay through it must be explicitly identified in item 11.2.2. Failing to do this would prevent legitimate relaying while simply allowing all relaying will lead to abuse of the relay.

11.2.4 Smart Host

This control determines whether this virtual server routes its outbound SMTP messages through a Smart host or whether it uses DNS to route its outbound messages.

Smart hosts can help secure communication, but configuring the entire virtual server to use the same smart host can lead to problems. As such, it is recommended that administrators NOT use smart hosts at the virtual server level. Instead, configure smart hosts on the SMTP connector as outlined in item 9.2.2.

Potential issues if this recommendation is not followed and smart hosts are used at the virtual server level:

- If there is more than one computer running Exchange in your organization, mail flow between servers may not work.
- If an IP address is listed in the Smart host text box and not enclosed in square brackets, mail flow may not work.
- If a name is listed in the Smart host box, it must be a Fully Qualified Domain Name (FQDN) or mail flow between servers may not work.

For additional information, refer to the Microsoft Document entitled: “Exchange Server 2003 Transport and Routing” located at http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/extransrout.mspx

11.2.5 Perform Reverse DNS Lookup on Incoming Messages

This feature causes the server to use a Directory Naming Service (DNS) lookup to try to determine the source of incoming email at the cost of a small amount of additional processor use per anonymous message.

While enabling this feature does not pose a security hazard, it is recommended that this feature be disabled as it is a waste of resources. It is relatively easy to fool the DNS lookup and so any results returned would be highly suspect. Given the lack of reliable information provided by this feature, it is usually not worth the performance cost to enable it.

11.2.6 TCP Port for Outbound Connections

This controls the port that the SMTP server uses when sending mail to other servers. The standard port for regular SMTP connections is 25. This port should be set to whatever port is used by all of the SMTP servers to which this SMTP server will be connecting.

Item 11.2.7 describes the field used to control the port to which a given SMTP Virtual Server will bind (listen). This control (item 11.2.6) controls the other side of this communication in that it controls the port on which this server sends (talks). If all the SMTP servers to which this server will be connecting will use a non-standard port, than that port must be entered here to allow this server to connect to them. However, since this field is used for ALL outbound communications, Exchange cannot be configured to use one port for some specified SMTP server, and another port for others. As such, unless every server which this server should connect to use the same port, this server will be unable to connect to all of these other servers. This is another reason (elaborated on more in item 11.2.7) why it may not be possible to change the SMTP port used in the Exchange organization.

Set this field to the port to which all connected SMTP are bound. Failure to do so will result in an inability to communicate with some or all other SMTP Virtual Servers.

11.2.7 TCP Port for SMTP Virtual Server

This controls the ports that the SMTP Virtual Server binds to for regular communications. The standard port for regular SMTP connections is 25. If ports other than these are to be used, clients will need to be explicitly configured to use the non-standard port.

Use the standard port (that is, the default). Traffic on these ports should be carefully monitored for signs of suspicious activity.

Changing the port to a non-standard value can provide some limited protection against automated attacks since these attacks will not connect to the correct port, and therefore be nullified. However, making this modification introduces a large amount of complexity for the system administrator: If clients are to connect to the server, they will all need to be configured to use the selected non-standard ports for their communication. In addition, if at some later point, additional network services are added to the Exchange server, care must be taken that the ports used by these services do not conflict with the non-standard ports chosen here. If clients outside the firewall or proxy server are to be allowed to connect to the server, or if automated network monitoring tools are employed, these services will all need to be configured to use the non-standard ports for SMTP traffic. Finally, most SMTP servers must connect to external SMTP servers when delivering and receiving mail. Unless all these other servers are configured to use the same non-standard port, this SMTP Virtual Server will be unable to connect to them. This means that an SMTP server that uses a non-standard port will probably be unable to send or receive mail from outside the organization. Even if ports are changed, however, a determined attacker may still be able to determine which port is used for the SMTP protocol by performing a comprehensive port scan, although doing so should be detectable by most network monitoring tools. Since changing the port introduces a large amount of complexity for a relatively small gain, the standard ports should be used.

11.2.8 Outbound Delivery Retry (group)

This controls the rate at which delivery attempts are retried as well as at what point the user is informed that their message is being delayed and at what point the server gives up on the delivery attempt. The control allows administrators to specify the “First retry interval” (amount of time before a second attempt is made to deliver a message), “Second retry interval” (time before a third attempt), “Third retry interval” (time before a fourth attempt), and “Subsequent retry interval” (time between all subsequent attempts). In addition to these fields, “Delay notification” field specifies the amount of time that must pass before a message is dispatched to the sender indicating that delivery of their message has not yet succeeded. Finally, the “Expiration timeout” gives the amount of time after which the server will give up attempting to deliver the message and simply discard it. Note that these controls only apply to messages sent from the domain.

The default values of these fields (10 minutes for the first three retries, 15 minutes for all subsequent retries, 12 hours before the user is notified of non-delivery, and 2 days before delivery is abandoned) should be adequate for most environments. Other factors such as the speed and reliability of the network and the capacity of the server may influence the ideal value of these settings. Administrators should monitor their system and may wish to modify the values as a result.

If the redelivery rate is too high (that is, if the time between redelivery attempts it too brief) then the server will spend extra time sending out repeated delivery attempts that have little chance of success. If the redelivery rate is too low, then messages will remain queued on the server longer than necessary. If one’s internal network is unreliable, then one may wish to increase the amount of time before a message is dropped after repeated delivery failures, since those failures are more likely to be local.

11.2.9 Maximum Hop Count

This controls the maximum number of steps (hops) which a message may take before it reaches its destination. The transmission of the message from one server to another is defined as a hop. This means that, for a hop count of 30, a message could pass through up to 29 intermediate servers before reaching its final destination. Messages that exceed their hop count are discarded undelivered. Hop counts were part of the original Internet protocol implementation to prevent the possibility of a routing glitch causing a message to be sent between machines indefinitely.

Recent studies indicate that virtually all messages can be delivered in under 25 hops. As such, the default hop count of 30 should suffice for most environments. Certain countries (such as France) tend to use a smaller number of fast gateways, which can increase the needed hop count (although the average hop count remains under 24). As such, enterprises located in, or which communicate extensively with certain countries may wish to increase their hop counts.[3] Administrators should monitor delivery – if messages to a particular destination or set of destinations are failing to be delivered on a regular basis, increasing the hop count may solve the problem. (It is more likely that the destination is experiencing server problems, but if this does not appear to be the case, the hop count is a valid possibility.)

If the hop count is set too low, messages are likely to expire before they can reach their destination. A hop count that is set too high is unlikely to have any noticeable effect, although in rare instances it may result in an undeliverable message cycling between intermediate servers for some period of time and contributing to network congestion.

11.2.10 Limit Number of Connections

This controls the maximum number of simultaneous outbound connections allowed to this SMTP Virtual Server. This can be used to throttle the use of the SMTP service if it begins to monopolize system resources. If the checkbox is cleared, no limits are applied to the number of connections. If the checkbox is selected, the text field on the right is used to specify the maximum number of connections. Note that this only limits the outbound connections – inbound connections are unaffected. To limit inbound connections, see item 11.2.13.

By default, the number of simultaneous outbound connections is limited to 1000. This should be sufficient for most enterprises. Exchange servers in especially large enterprises with lots of users, or servers which otherwise experience a large amount of load as part of their daily use may wish to increase this amount. Administrators should monitor SMTP usage and may wish to modify this value to better reflect enterprise requirements.

If the limit on connections is too low, connections are likely to be dropped resulting in delays for users wishing to send messages. On the other hand, if the limit is too high, it may result in a disproportionately large amount of server resources being employed by the SMTP server. (That said, SMTP connections are not particularly costly and are likely to be mission critical in the average Exchange environment. As such, administrators should probably increase the number of connections more readily than decrease it.)

11.2.11 Time Out (minutes)

This controls the number of minutes an idle connection will be maintained by the server before it is dropped. This can be used to help limit the number of simultaneous connections the server must support. The speed with which idle connections are dropped is particularly important if the number of simultaneous connections permitted (using item 11.2.10) is low, since idle connections may take up a large number of the available connections. Note that this only affects outbound connections. To set the timeout for idle inbound connections, see item

[3] "Measurements On Delay And Hop-Count Of The Internet", Aiguo Fei, Guangyu Pei, Roy Liu, and Lixia Zhang, IEEE GLOBECOM'98 - Internet Mini-Conference, 1998

11.2.14.

It is recommended that this value be set to 10 minutes. This provides users with a reasonable window in which to resume activities without maintaining idle connections for excessive intervals. Administrators should monitor usage and modify this value as necessary. If administrators find that they are frequently hitting the limit on the number of simultaneous connections provided in item 11.2.10 (but they feel this connection limit is reasonable for their environment) then the number of minutes until timeout should be reduced.

If connections time out after too few minutes then connections may be terminated unexpectedly. Although most clients and servers will transparently reconnect if additional messages need to be sent, there is some network overhead to (re)establishing a connection. As such, it is in the interest of both parties to maintain connections if it is likely that further communication will take place in the near future. On the other hand, if the timeout period is too long, large numbers of unused connections may end up being maintained for an unnecessarily long period of time. This is unlikely to affect resources, since idle connections take up little processing power or bandwidth, but may result in the number of open connections hitting the limit prescribed by control item 11.2.10, which could prevent new connections from being established.

11.2.12 Limit Number of Connections Per Domain

This controls the maximum number of simultaneous outbound connections allowed to the SMTP server from any single internal domain. This can be used to throttle the use of the SMTP service if some domains end up taking up a larger percentage of the total outbound SMTP connections (item 11.2.10). If the checkbox is cleared, no limits are applied to the number of connections. If the checkbox is selected, the text field on the right is used to specify the maximum number of connections. Note that this only limits the outbound connections – inbound connections are unaffected.

By default, the number of simultaneous outbound connections from a single domain is limited to 100. This should be sufficient for most enterprises. Exchange servers that service a smaller number of large domains may wish to increase this number to give each domain a larger proportion of the available outbound connections (item 11.2.10). Administrators should monitor SMTP usage and may wish to modify this value to better reflect enterprise requirements. For example, if different domains use this server in roughly equal proportion, the administrator may wish to change the setting to be closer to the total number of available SMTP connections divided by the number of domains (thus enforcing this equality) or possibly even disable this limit entirely by clearing the checkbox. If, however, certain domains are expected to account for a unusually large number of the available SMTP connections, administrators may wish to set the limit to a value higher than what would be used for an equal distribution. Doing this would allow the high usage domains to continue to operate at their expected levels but also ensure that they do not completely shut the other domains out of the server.

If the limit on connections is too low for any particular domain, connections are likely to be dropped resulting in delays for users wishing to send messages. On the other hand, if the limit is too high, it may result in a few domains taking up a disproportionately large amount of the available SMTP connections and possibly denying access to other domains.

11.2.13 Limit Number of Simultaneous Connections to Virtual Server

This controls the maximum number of simultaneous inbound connections allowed to the SMTP server. This can be used to throttle the use of the SMTP service if it begins to monopolize system resources. If the checkbox is cleared, no limits are applied to the number of connections. If the checkbox is selected, the text field on the right is used to specify the maximum number of connections. Note that this only limits the inbound connections – outbound connections are unaffected. To limit outbound connections, see item 11.2.10.

By default, the number of simultaneous inbound connections is unlimited. In most enterprises, SMTP will be a central service of Microsoft Exchange. As such, administrators should be reluctant to limit its accessibility. On the other hand, if other services are intended to share resources on this server with SMTP, it might be reasonable to apply some limit to inbound connections in order to prevent heavy use of SMTP from monopolizing resources and preventing proper functioning of other services. Administrators should monitor activity to determine if a limit is necessary and what that limit should be.

If the number of SMTP connections is limited, external clients will need to share from a limited pool of connections. If the limit is high enough, this will not be a problem, but if the limit is too low, the pool of available connections may get filled causing attempts to create new connections to fail and resulting in mail delivery failure. (Most clients will attempt to redeliver failed messages, but at the very least, this will delay message delivery.) Moreover, if a limit is imposed then external attackers only need to initiate and maintain that many simultaneous connections in order to deny service to the SMTP server. On the other hand, if the number of SMTP connections is unlimited, attackers will have a much more difficult time disabling SMTP since they will need to maintain so many connections that the server has trouble finding resources to maintain them all. (Of course, if an attacker is able to do this, and SMTP is unlimited, it means that such a denial of service attack would block all services, not just SMTP.)

11.2.14 Connection Time-out (minutes)

This controls the number of minutes an idle connection will be maintained by the server before it is dropped. This can be used to help limit the number of simultaneous connections the server must support. The speed with which idle connections are dropped is particularly important if the number of simultaneous connections permitted (using item 11.2.13) is low, since idle connections may take up a large number of the available connections. Note that this only affects inbound connections. To set the timeout for idle outbound connections see item 11.2.11.

It is recommended that this value be set to 10 minutes. This provides users with a reasonable window in which to resume activities without maintaining idle connections for excessive intervals. Administrators should monitor usage and modify this value as necessary. If administrators find that they are frequently hitting the limit on the number of simultaneous connections provided in item 11.2.13 (but they feel this connection limit is reasonable for their environment) then the number of minutes until timeout should be reduced.

If connections time out after too few minutes then connections may be terminated unexpectedly. Although most clients and servers will transparently reconnect if additional messages need to be sent, there is some network overhead to (re)establishing a connection. As such, it is in the interest of both parties to maintain connections if it is likely that further communication will take place in the near future. On the other hand, if the timeout period is too long, large number of unused connections may end up being maintained for an unnecessarily long period of time. This is unlikely to affect resources, since idle connections take up little processing power or bandwidth, but may result in the number of open connections hitting the limit prescribed by item 11.2.13, which could prevent new connections from being established.

11.3 Delivery Restrictions

11.3.1 Exclude or Limit Connections

This controls which IP addresses are allowed to connect to this virtual server to download messages. The control can be set to either allow all computers to connect except for a specified few, or to deny all computers except for a specified few. In addition to individual IP addresses, subnet masks may be used to specify groups of addresses. Note that if your network uses DHCP to dynamically assign IP addresses to client computers it will not be possible to reliably specify one computer out of the group serviced by a given DHCP server since that computer's IP address may change regularly. (If all computers serviced by a DHCP server should have the same treatment then there is no problem since one or more subnet masks can then be used to cover the DHCP server's entire address pool, ensuring that regardless of the address a computer is assigned, it will have the same treatment in this panel.) For additional information please refer to the Exchange 2003 Transport and Routing Guide http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/extransrout.mspx and/or the What's New in Exchange Server 2003 http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/newex03.mspx

Select “Only the list below” so that the administrator must explicitly specify which clients and other Exchange servers can connect to the SMTP Virtual Server. This significantly reduces the chance of unauthorized connections to the server and helps to further restrict network connectivity.

If “All except the list below” must be selected, administrators should monitor connectivity to the SMTP server to ensure that no suspicious connections are being made

11.3.2 Apply Sender/Recipient/Connection Filters

These checkboxes are used to control which filters are applied to messages passing through this virtual server. The configuration of the sender filter is described in items 1.2.8 - 1.2.12, the configuration of the recipient filter is described in items 1.2.6 - 1.2.7, and the configuration of the connection filter is described in items 1.2.1 - 1.2.3. Only when the respective checkboxes are selected here will the corresponding filters be applied to SMTP traffic.

Select all three of these checkboxes. (Even if the actual filter elements in the Message Delivery properties panel are not actually configured to perform any action, it is still recommended to select these checkboxes since it reduces the chance that someone at some later point might choose to configure a filter, but forgets to select this checkbox and, as a result, fail to have the filter take effect.) The justification for the individual filter elements are discussed in the sections for the relevant controls.

Failure to select a checkbox will result in the corresponding filter not being applied.

11.3.3 Limit Message Size

This controls the maximum permitted size of messages sent to or from this virtual server. If the checkbox is cleared, no restrictions on message size are applied by the virtual server. If the checkbox is selected, the field on the right of the panel is used to specify the maximum message size. This limit applies not only to user mail but also to other features that use SMTP, such as public folders. Limits may also be applied across connectors (items 9.3.5 and 9.3.7), to e-mail messages in particular (item 1.2.4), and per sender/recipient (items 8.2.1 and 8.2.2 for folders, and item 12.2.3 for users). When multiple limits apply, the most restrictive limit for the appropriate message applies. For example, if a message is sent to a particular user, and limits are applied for that user at the virtual server, and applied to e-mail messages overall, the message must be allowed by all three limits in order to reach the user.

The default setting of disabling size limits at the virtual server is recommended. Limits on e-mail messages are already recommended in item 1.2.4 which, if followed, should provide sufficient protection against excessively large messages passing through the virtual server. Applying limits to email messages rather than at the virtual server level allows limits to be applied to messages without unnecessarily restricting the use of SMTP by other features in Exchange.

If the limit is set too low, legitimate messages may be rejected if they exceed the size limits. This may prevent messages from reaching recipients or prevent certain features of Exchange from functioning. If, however, no limits are applied to messages either here or through the other mechanisms outlined above, excessively large messages may be able to consume large amounts of server resources. As such, some limits should be applied, although the virtual server may not be the best place to do so.

11.3.4 Limit Session Size

This controls the maximum permitted size of SMTP sessions (inbound and outbound) on this virtual server. If the checkbox is cleared, no restrictions on session size are applied by the virtual server. If the checkbox is selected, the field on the right of the panel is used to specify the maximum session size. This limit applies to all use of the SMTP protocol, including the transmission of email as well as other features, such as Public Folders.

By default, no limits on session size are applied and it is recommended that this setting be preserved. If a size limit was imposed a large session would simply be broken up into multiple smaller sessions resulting in additional overhead. If limits are applied to messages, that should provide sufficient protection against excessively large communications. Sometimes multiple sessions can increase efficiency by allowing sessions to be processed in parallel, but this is best done by the number of messages in a session (as controlled by item 11.3.5) rather than the size of the session itself. As such, it is recommended that no limit be applied.

If the limit is set too low, the SMTP server will be forced to process a larger number of smaller sessions in order to relay the same number of messages. If, on the other hand, one begins to observe SMTP sessions of excessive size which are not being stopped by message filters (that is, one is receiving sessions consisting of a very large number of small messages) then limiting session sizes may be an option. However, this is an unlikely occurrence.

11.3.5 Limit Number of Messages Per Connection

This controls the maximum number of messages allowed in a single SMTP session. By limiting the number of messages in a session, delivery efficiency can be improved by allowing the processing of multiple sessions in parallel. If the checkbox is cleared, no limit is set on the number of messages in a single session. If the checkbox is selected, the field on the right of the panel is used to specify the message limit.

This control does not affect security so there is no security risk associated with any particular setting. Instead, this control attempts to improve performance by breaking large numbers of messages into multiple sessions. No messages will be prevented from being sent by this control. If the message limit is set to too high a value (or turned off completely) one is unlikely to reap many benefits from the limit since it will increase the chance that a batch of messages will fit in a single session and hence cannot be processed in parallel. On the other hand, if the limit is set too low, the advantages in processing multiple sets of messages in parallel will be outweighed by the network cost of establishing a large number of sessions. In general, the default value of 20 will be appropriate for most environments.

11.3.6 Limit Number of Recipients per Message

This field is used to control the maximum number of recipients that can be specified in a single message sent from this server. It is intended to improve efficiency by forcing messages sent to a greater number of recipients to be sent out in multiple messages. Note that this is not the same as item 1.2.5 which attempts to prevent mass mailing by imposing a maximum number of recipients of any message. This control, on the other hand, simply attempts to improve efficiency by parallelizing the transmission of a message with a very large recipient list. Item 1.2.5 rejects any message with too many recipients; this feature simply causes the message to be copied with each copy sent to fewer than the specified number of recipients. If the checkbox is cleared, no limits are applied. If the checkbox is selected, the field to the right is used to hold the maximum number of recipients.

This control does not affect security so there is no security risk associated with any particular setting. Instead, this control attempts to improve performance by breaking large numbers of recipients into multiple messages. No messages will be prevented from being sent by this control. If the number is too high, the server will miss out on the efficiency increase that comes from parallelizing the message transmission. On the other hand, if the number of recipients is too low, messages may be copied more often then necessary which will result in a slight additional load on the server. In general, the default value of 64000 will be appropriate for most environments

11.4 logging

11.4.1 Enable Logging

This controls the creation and format of log files used to monitor the interaction between this SMTP Virtual Server and other SMTP hosts. By default, these files will be stored in \WINNT\SYSTEM32\LOGFILES\SMPTVSx (where x is a number used to distinguish between virtual servers in this organization). The drop-down menu is used to select the format of the log file. The properties button next to this dropdown displays configuration information specific to the type of log format selected, but usually has some control to indicate the log rotation schedule (that is, how often the old log file should be closed and a new log file should be started).

Logging should be enabled. Doing so records attempts made to connect to this virtual server. In the case of an attack on the SMTP server, these logs could contain useful details regarding the time and nature of the attack. Due to the size of log files, the files should be regularly copied to external storage and deleted from the server to conserve memory. Log files should be retained for at least one month. The format of the log files is largely a matter of preference for the administrator.

Failure to enable this feature means that the administrator will not have access to information regarding SMTP connections should the SMTP servers come under attack.

It is required that all log files be written to separate partitions from the partitions used by the Exchange Stores. Exchange will dismount its stores if it detects that it has run out of disk space, resulting in a complete loss of Exchange services. To minimize the chance of this happening, log files should write to some other partition so that if the logs fill this partition it will not result in the failure of Exchange.

12. User Level Controls

12.1 Deletion Settings

12.1.1 Deleted Item Retention (group)

This control allows the mailbox store's deletion settings (items 4.3.1 and 4.3.2) to be overridden for this user. If the "Use mailbox store defaults" checkbox is selected, the deletion settings set in the mailbox store will be applied for this user. If the checkbox is cleared, then the settings in the mailbox store will not be applied and the "Keep deleted items for (days)" and "Do not permanently delete items until this store has been backed up" fields in this group will be used to control deletion item retention for this user. These correspond to items 4.3.2 and 4.3.1, respectively. See these controls for recommendations and risks involved with the configuration of these items. Note that if the mailbox store defaults are overridden then both items 4.3.1 and 4.3.2 will be overridden by the settings in the associated fields in this panel. It is not possible to just override one. As such, if the mailbox store defaults are overridden, be sure that both fields reflect the desired settings.

It is recommended that the mailbox store defaults be used. There may be cases where it is reasonable or necessary to change the settings of a particular user, but these cases should be limited as much as possible.

If a user's configuration overrides the mailbox store defaults then that configuration needs to be administered independently. This means that, when operational or capacity requirements change, in addition to updating the mailbox store, the settings of all the independently administered users must be updated as well. This increases the chance that the user's configuration may be overlooked and end up configured in a way that is no longer appropriate. While there are certainly circumstances where overriding the mailbox store settings is the correct course of action, doing so does require additional time and attention on the part of the administrator.

12.2 Delivery Restrictions

12.2.1 Send on Behalf

This controls whether other users can send mail on behalf of this user. For example, if this user's email name is "Adam" and the user "Bob" is added to this field in Adam's Delivery Options properties, then Bob will be able to send mail on behalf of Adam. The from field of the resulting message will look like:

From: Bob on behalf of Adam

It is recommended that the "Send on behalf" field be left blank (messages are not sent on behalf of any party). While the full from field displays both the actual sender as well as who the message is on behalf of, in many instances only the party on whose behalf the message was sent may be seen. For example, if the message is viewed only in the preview pane of Outlook, the actual sender may never be seen. As such, it is possible for a message's actual sender to be effectively hidden from the user using this feature.

If users are allowed to send on behalf of other parties, then users may never realize the actual sender of the message. This can allow senders to mask their activities. As such, if Send on behalf is used it is important that only trusted parties are given the ability to send on behalf of others, and their messages should be reviewed to ensure this privilege is not being abused.

12.2.2 Recipient Limits

This control allows the settings of item 1.2.5 (Global recipient limits) to be overridden for this user. If “Use default limits” is selected, the recipient limits set in item 1.2.5 will apply to this user. If “Maximum recipients” is selected, the default set in item 1.2.5 will be overridden and the field to the right of the panel will specify the recipient limits that will apply to this user (the maximum number of recipients that this user will be able to send to in a single message). See item 1.2.5 for policy recommendations and risks regarding recipient limits.

It is recommended that, in general, users use the setting prescribed in the global settings. There may be cases where it is reasonable or necessary to change the settings of a particular user, but these cases should be limited as much as possible.

If a user’s configuration overrides the global settings then that configuration needs to be administered independently. This means that, when operational or capacity requirements change, in addition to updating the global settings, the settings of all the independently administered users must be updated as well. This increases the chance that the user’s configuration may be overlooked and end up configured in a way that is no longer appropriate. While there are certainly circumstances where overriding the global settings is the correct course of action, doing so does require additional time and attention on the part of the administrator.

12.2.3 Sending Message Size/Receiving Message Size

These control whether this user will use the default limits (set in item 1.2.4) or whether the user will observe different limits. Both the sending message size limits and the receiving messages size limits can be overridden independently, but their controls function in the same way. If “Use default limit” is selected, the user will observe the sending/receiving message size limits set in the global policy in item 1.2.4. If “Maximum KB” is selected, the sending/receiving message size limits as set in the global settings will be overridden and the text field to the right of “Maximum KB” will be used to hold the user’s personalized sending/receiving size limit. See item 1.2.4 for recommendations regarding sending and receiving message size limits and the risks of various settings.

It is recommended that, in general, users use the setting prescribed in the global settings. There may be cases where it is reasonable or necessary to change the settings of a particular user, but these cases should be limited as much as possible.

If a user’s configuration overrides the global settings then that configuration needs to be administered independently. This means that, when operational or capacity requirements change, in addition to updating the global settings, the settings of all the independently administered users must be updated as well. This increases the chance that the user’s configuration may be overlooked and end up configured in a way that is no longer appropriate. While there are certainly circumstances where overriding the global settings is the correct course of action, doing so does require additional time and attention on the part of the administrator.

12.2.4 Message Restrictions (group)

This allows administrators to indicate that this user should not receive mail from certain senders. It is used in addition to the limits imposed by items 1.2.2 (blacklist provider intended to block spammers) and 1.2.8 (intended to block certain senders on an enterprise wide basis). This means that, even if this control specifies that the user is allowed to receive mail "From everyone", mail from senders given in items 1.2.2 or 1.2.8 will still be blocked. This field is simply intended to block mail from senders who should not be sending mail to this user in particular. This control allows mail to be allowed "From everyone", from a limited set of users, or from everyone except a limited set of users. In addition, a further restriction may be applied to any of these settings in that only authenticated users may be allowed to send to this user.

It is recommended to select the "From everyone" checkbox, thus providing no additional limits on who may send mail to this user. In most cases where a sender should be blocked, the global level filters described in items 1.2.2 and 1.2.8 would be the correct place to specify the denied user. For most users, limiting their mail to come only from a finite set of users, or only users who can authenticate to Exchange would be too restrictive. In limited circumstances, it might be necessary for a user to be excluded from receiving mail from a small set of users, but not to block those users at the domain. However, such a case would be very rare. Moreover, it would be of dubious utility since it is relatively easy for any of these illicit senders to create a new email address which would not be blocked. As such, "From everyone" is the most appropriate setting for most cases. (Note that, for groups, we recommended limiting senders to authenticated users to try to protect distribution groups from external spammers (item 12.2.5). While this setting makes sense for groups, it is likely too restrictive for individual users.)

If a user is barred from receiving mail from unauthenticated users or from all users except a small set, this means that the user will be unable to use email to communicate with anyone except for the limited party specifically allowed. This is likely to limit a great deal of potentially reasonable interaction. Moreover, if a new, legitimate user appears, the user's configuration may need to be manually changed to allow this user to communicate with them.

12.2.5 Accept Messages: From Authenticated Users Only

This item controls whether the specified mail enabled group will only be allowed to receive messages from authenticated senders. Since senders can only be authenticated if they are members of the group's domain, this effectively limits senders to entities within the organization, blocking messages from external sources.

This setting can often be employed to great effect in distribution groups. Spammers often target distribution groups since it results in a large number of recipients for each message sent out. If only authenticated messages to this group are permitted, this means that spammers will be unable to use this group to spread their messages. If a group does not have a legitimate reason to pass on mail from external sources, enable this feature. Groups which should be recipients of external mail, however, may not be able to operate with this limitation. As such, there may be cases where this recommendation would not apply.

Note that a similar feature can be enabled for users to limit senders of messages to authenticated users (item 12.2.4). While the feature might be useful in limited circumstances (for example, if sending to a help desk address which has no reason to receive external email) for most users, the inability to receive external email is too restrictive for daily use.

12.3 Enabled Services

12.3.1 Enable POP3

This controls whether this user is permitted to use the POP3 service. This can be used to deny a user the ability to use this protocol even if the protocol is enabled on the server. (Obviously, if the protocol is disabled on the server, the user will not be able to use it regardless of the setting of this control.) To change the setting for this protocol, select it from the list and then click the "Enable" or "Disable" button below the list as appropriate. It is also possible to customize some of the settings of this protocol using this interface, but these settings are largely formatting issues and none are security relevant so they will not be discussed further here.

Note that any risks presented by POP3 will be exposed when the service is enabled regardless of whether any internal users can use the protocol. As such, disabling the use of POP3 at the user level does almost nothing to enhance system security. There may be isolated cases where certain users should be prevented from using POP3 for operational reasons – that is, a certain user should not be allowed to use the protocol for policy reasons, although such cases are likely to be rare. As such, no recommendation is made.

While disabling (or enabling) a user's access to the protocol does not significantly affect the security of the system, it does impact the ease with which the environment can be administered. Users for whom non-default settings are given must be carefully tracked to ensure that, in the case of a global policy change, the user's customized configurations continue to comply with overall stances.

12.3.2 Enable IMAP

This controls whether this user is permitted to use the IMAP service. This can be used to deny a user the ability to use this protocol even if the protocol is enabled on the server. (Obviously, if the protocol is disabled on the server, the user will not be able to use it regardless of the setting of this control.) To change the setting for this protocol, select it from the list and then click the "Enable" or "Disable" button below the list as appropriate. It is also possible to customize some of the settings of this protocol using this interface, but these settings are largely formatting issues and none are security relevant so they will not be discussed further here.

Note that any risks presented by IMAP will be exposed when the service is enabled regardless of whether any internal users can use the protocol. As such, disabling the use of IMAP at the user level does almost nothing to enhance system security. There may be isolated cases where certain users should be prevented from using IMAP for operational reasons – that is, a certain user should not be allowed to use the protocol for policy reasons, although such cases are likely to be rare. As such, no recommendation is made.

While disabling (or enabling) a user's access to the protocol does not significantly affect the security of the system, it does impact the ease with which the environment can be administered. Users for whom non-default settings are given must be carefully tracked to ensure that, in the case of a global policy change, the user's customized configurations continue to comply with overall stances.

12.3.3 Enable OWA

This controls whether this user is permitted to use the Outlook Web Access (OWA) service. This can be used to deny a user the ability to use this service even if the protocol is enabled on the server. (Obviously, if the service is disabled on the server, the user will not be able to use it regardless of the setting of this control.) To change the setting for this service, select it from the list and then click the “Enable” or “Disable” button below the list as appropriate.

Note that any risks presented by OWA will be exposed when the service is enabled regardless of whether any internal users can use the service. As such, disabling the use of OWA at the user level does almost nothing to enhance system security. There may be isolated cases where certain users should be prevented from using OWA for operational reasons – that is, a certain user should not be allowed to use the protocol for policy reasons, although such cases are likely to be rare. As such, no recommendation is made.

While disabling (or enabling) a user’s access to the service does not significantly affect the security of the system, it does impact the ease with which the environment can be administered. Users for whom non-default settings are given must be carefully tracked to ensure that, in the case of a global policy change, the user’s customized configurations continue to comply with overall stances.

12.3.4 Enable OMA

This controls whether this user is permitted to use the Outlook Mobile Access (OMA) service. This can be used to deny a user the ability to use this service even if the protocol is enabled on the server. (Obviously, if the service is disabled on the server, the user will not be able to use it regardless of the setting of this control.) To change the setting for this service, select it from the list and then click the “Enable” or “Disable” button below the list as appropriate.

Note that any risks presented by OMA will be exposed when the service is enabled regardless of whether any internal users can use the service. As such, disabling the use of OMA at the user level does almost nothing to enhance system security. There may be isolated cases where certain users should be prevented from using OMA for operational reasons – that is, a certain user should not be allowed to use the protocol for policy reasons, although such cases are likely to be rare. As such, no recommendation is made.

While disabling (or enabling) a user’s access to the service does not significantly affect the security of the system, it does impact the ease with which the environment can be administered. Users for whom non-default settings are given must be carefully tracked to ensure that, in the case of a global policy change, the user’s customized configurations continue to comply with overall stances.

12.3.5 Enable ActiveSync

This controls whether this user is permitted to use the ActiveSync service. This can be used to deny a user the ability to use this service even if the protocol is enabled on the server. (Obviously, if the service is disabled on the server, the user will not be able to use it regardless of the setting of this control.) To change the setting for this service, select it from the list and then click the “Enable” or “Disable” button below the list as appropriate.

Note that any risks presented by ActiveSync will be exposed when the service is enabled regardless of whether any internal users can use the service. As such, disabling the use of ActiveSync at the user level does almost nothing to enhance system security. There may be isolated cases where certain users should be prevented from using ActiveSync for operational reasons – that is, a certain user should not be allowed to use the protocol for policy reasons, although such cases are likely to be rare. As such, no recommendation is made.

While disabling (or enabling) a user’s access to the service does not significantly affect the security of the system, it does impact the ease with which the environment can be administered. Users for whom non-default settings are given must be carefully tracked to ensure that, in the case of a global policy change, the user’s customized configurations continue to comply with overall stances.

12.3.6 Up-to-Date Notifications

This controls whether this user is permitted to use the up-to-date notifications feature of the ActiveSync service. This can be used to deny a user the ability to use this service even if the protocol is enabled on the server. (Obviously, if the service is disabled on the server, the user will not be able to use it regardless of the setting of this control.) To change the setting for this service, select it from the list and then click the “Enable” or “Disable” button below the list as appropriate. Note that up-to-date notifications for ActiveSync can only be enabled if user initiated synchronization (also known as, ActiveSync) is also enabled.

Note that any risks presented by up-to-date notifications for ActiveSync will be exposed when the service is enabled regardless of whether any internal users can use the service. As such, disabling the use of up-to-date notifications for ActiveSync at the user level does almost nothing to enhance system security. There may be isolated cases where certain users should be prevented from using up-to-date notifications for ActiveSync for operational reasons – that is, a certain user should not be allowed to use the protocol for policy reasons, although such cases are likely to be rare. As such, no recommendation is made.

While disabling (or enabling) a user’s access to the service does not significantly affect the security of the system, it does impact the ease with which the environment can be administered. Users for whom non-default settings are given must be carefully tracked to ensure that, in the case of a global policy change, the user’s customized configurations continue to comply with overall stances.

12.4 Permissions

12.4.1 Mailbox Rights

This controls simple permissions for access to a user's mailbox. This allows administrators to assign the ability to read or delete messages.

The default configuration is appropriate in virtually all cases. There may be rare instances when it is necessary to give additional users access to a user's mailbox, however, doing this should be done only rarely.

If additional users are given access to a user's mailbox they will have special access with regards to this user's mail account. This represents a security risk so any such persons added to a user's account must be highly trusted. If a user's access to their own mailbox is reduced they may not be able to utilize normal mail functionality.

12.5 Storage Limits

12.5.1 Storage Limits (groups)

This control allows the settings of item 4.7.1 (mailbox storage limits) to be overridden for this user. Specifically, it allows the administrator to customize how much memory a user's mailbox may consume before they receive warning messages, before they are no longer able to send mail, and before they can neither send nor receive. If "Use mailbox store defaults" is selected, the storage limits set in item 4.7.1 will apply to this user. If this checkbox is cleared, the fields for "Issue warning at", "Prohibit send at", and "Prohibit send and receive at" will be enabled allowing customized values to be placed in these fields. See item 4.7.1 for policy recommendations and risks regarding recipient limits. Note that it is not possible to override just one of the three fields. If "Use mailbox store defaults" is not selected then all storage limit controls in the mailbox store will be ignored for this user. This means that, if the user overrides the mailbox store and fails to enable one of the three features, then that feature will not be enabled for this user regardless of the setting on the mailbox store. As such, be sure that the settings of all three fields reflect the desired actions if the mailbox store's configuration is overridden.

It is recommended that the mailbox store defaults be used. There may be cases where it is reasonable or necessary to change the settings of a particular user, but these cases should be limited as much as possible.

If a user's configuration overrides the mailbox store settings then that configuration needs to be administered independently. This means that, when operational or capacity requirements change, in addition to updating the mailbox store settings, settings of all the independently administered users must be updated as well. This increases the chance that the user's configuration may be overlooked and end up configured in a way that is no longer appropriate. While there are certainly circumstances where overriding the mailbox store settings is the correct course of action, doing so does require additional time and attention on the part of the administrator.

12.6 Visibility

12.6.1 Hide from Exchange Address List

This allows a user's email address to be hidden from Exchange's address lists. If selected, the user's name will not appear in any Exchange address lists.

Ideally, all users would be hidden from Exchange address lists. This would reduce the dissemination of email addresses and make it more difficult for spammers and other malicious third parties to learn internal email addresses. However, doing this would be extremely time-consuming and error prone since there is no mechanism to modify this setting except by manually changing the setting for each user individually. As such, it is acceptable to leave the setting at the default.

Leaving the setting at the default (cleared – address is not hidden from lists) means that the use of address lists should be tightly controlled. Moreover, spam filters should be put in place (items 1.2.1 and 1.2.2) to limit the amount of junk email sent to accounts. If addresses are hidden (the checkbox is selected) administrators will need to ensure that new users continue are also updated to comply with this policy.

12.6.2 ILS Server Settings

This is used to set the Internet Locator Service (ILS) server and account name for a user. ILS is used to dynamically keep track of which users are online at any given time. When users log on or off the ILS updates its list of active users. Other users can then use this list to determine if colleagues or friends are available online. The fields of this panel are used to set the ILS server that will record the liveness of this user, as well as the user's account name as it will appear in the ILS's list of active users. If these fields are left blank, the user will not register with any ILS server.

It is recommended that these fields be left blank and that ILS not be used. External ILS servers present a particular risk in that it broadcasts user information (such as the user's IP address). This can reveal information about system architecture and active user accounts that may be of use to attackers. Internal ILS servers represent a much smaller risk and may be used if closely monitored.

An external ILS server would collect a great deal of sensitive information about users and the enterprise and make this information available to other parties. Internal ILS servers pose a much smaller threat, although access to them and their data should be tightly controlled. Information stored by ILS servers, such as IP addresses of users, can be used by hostile parties to mount attacks on the enterprise.

This Page Intentionally Left Blank.

Appendix

A

Bibliography

Exchange Server 2003 Security Hardening Guide; Michael Grimm and Michael Nelte; The Microsoft Corporation; February 2004.

Exchange Server 2003 Deployment Guide, The Microsoft Corporation (http://www.microsoft.com/technet/prodtechnol/exchange/guides/Ex2k3DepGuide/f9918adf-057a-4235-8f7e-f7f27f3a8789.mspx)

Exchange Server 2003 and Exchange Server Front-End and Back-End Topology; Joey Masterson and Andrew Moss, The Microsoft Corporation, July 2004.

Mastering Microsoft Exchange Server 2003. Barry Gerber, SYBEX Inc., 2003.

Windows Server 2003 Security Guide; Kurt Dillard, Jose Maldonado, and Brad Warrender; The Microsoft Corporation, 2003.

Windows Server 2000 Operating System Level 2 Benchmark Consensus Baseline Security Settings; Jeff Shawgo, Editor; The Center for Internet Security; Version 1.02, 2 September 2003.