

Industry Insight

BNA INSIGHT

This two-part article is a comprehensive primer for legal and RIM professionals on aspects of Exchange 2010 relating to records management and eDiscovery.

Part I, published at 13 DDEE 159, explored several aspects of Exchange 2010: its history and evolution; deployment options; and, its “core” and “extended” footprints. Part II explores the native (out-of-the-box) records management and eDiscovery features of Exchange 2010.

Microsoft Exchange 2010: What Legal and RIM Professionals Should Know—Part II



By JOHN P. COLLINS

INTRODUCTION

The goal of this article is to examine how Exchange 2010's records management and eDiscovery features work, and identify issues which could arise from deployment and use of those features.

A number of interrelated elements comprise Exchange 2010's records management and eDiscovery feature-set. Ensuring these elements are configured properly and interoperate as planned is essential if the records management and eDiscovery processes supported by Exchange 2010 are to follow—or at least aspire to follow—best practices. For example:

- Although the new personal archive can store dozens of gigabytes of email, it should be implemented in tandem with PST file elimination.
- Retention tags can be used to specify when email can be purged, but a litigation hold policy and a process suspending such purging also must be in place (Exchange 2010 has a built-in litigation hold feature).

- While the discovery search feature enables performing a single search across all of an organization's mailboxes (using keywords and other parameters), certain file types are not, by default, searchable—certainly a critical consideration from an eDiscovery perspective.

This article examines these features—and others—in an effort to provide legal and RIM professionals with an understanding of whether—and how to—leverage the native records management and eDiscovery features of Exchange 2010.

Centralizing Storage of Email¹

One of the first considerations when implementing Exchange 2010's records management and eDiscovery

¹ While beyond the scope of this article, any discussion about centralizing the storage and management of email should also address backup media retention policies. The retention policies for backup and server based email should be carefully synchronized.

John P. Collins is the Founder and Vice President of Consulting for The Ingersoll Firm. The Ingersoll Firm is the nation's leading ESI data mapping consultancy, working with Fortune 500 companies, AMLAW 100/200 law firms, and electronic discovery services providers. Mr. Collins has created dozens of ESI data maps for a wide variety of companies in the pharmaceutical, manufacturing, health services, and consumer products industries. Mr. Collins has an Executive Juris Doctorate from Concord Law School and 16 years of information technology (IT) experience. He can be reached at jcollins@theingersollfirm.com.

features is centralizing the storage of email on the Exchange Server itself. Email outside the control of Exchange—for example, residing in PST files—is beyond the control and order imposed by Exchange 2010's records management and eDiscovery features.

Consequently, when contemplating deployment of Exchange 2010's records management and eDiscovery features, it is essential to create a plan to bring all email into the Exchange environment proper.²

To support such a plan, Exchange 2010 introduced two new features: large mailboxes and personal archives.

Large Mailbox Support

Large mailboxes are not a feature per se, but rather, an architectural change in the underlying Exchange 2010 database and storage system. Prior to Exchange 2010, the architecture and storage was optimized for small mailboxes. Larger mailboxes introduced exorbitant system processing overhead, extended backup window requirements, and increased storage costs. Exchange 2010 addresses these issues through a re-architected storage system that embraces the concept of large (multi-gigabyte) mailboxes.

Large mailbox support means organizations no longer need to employ mailbox quotas to keep mailboxes below a certain size. In the absence of quotas, users don't have to move items into PST files to avoid exceeding their quota.

While current users could, in theory, amass enormous volumes of email (20 or 30 gigabytes for example), as discussed in this article, Exchange 2010 provides records management and archiving features that support implementing an email management policy that calls for systematic purging of email.

Large mailbox support is a key foundational element for eliminating PST files, and thus enabling the centralized storage of email.

Personal Archive

Exchange 2010 introduces the personal archive, an additional mailbox designed to augment a user's primary mailbox and replace the PST files into which users have historically offloaded email messages. Personal archives, along with large mailbox support, enable bringing large volumes of (formerly unmanaged) email under control.

By eliminating PST files, challenges familiar to many legal and RIM professionals are addressed, such as:

- Where are PST files located? (Removable media, file shares, local hard drive, etc.?)
- Have all PST files been identified?
- Is data or information potentially responsive to a discovery request contained within unidentified PST files?
- Is it feasible to apply retention policies or litigation holds against email residing in PST files?

² Ideally, PST files would be eliminated, and users would not be permitted to store email outside of Exchange, except in a managed repository (for example, an email archiving solution). There are a number of approaches to centralizing storage of email within the Exchange environment proper. If users are allowed to store email outside of the Exchange environment, then such locations should integrate with records management and eDiscovery tools.

The personal archive helps solve these (and other) challenges by providing an Exchange-based storage location for email that previously would likely have been stored in a PST file. The personal archive, as a component of the Exchange mail store, is subject to all the management features and functions of Exchange 2010—including the records management and eDiscovery feature set.

Key aspects of personal archives include:

- Each user may have a single personal archive (it must be activated);
- A default retention policy (which can be changed or modified) is automatically applied to each user with a personal archive;
- The personal archive is indexed and searchable for discovery purposes;
- All items residing in the personal archive are subject to litigation holds;
- PST files can be moved into the personal archive (and the folder structure of the PST file is retained); and
- Users can search the personal archive and their primary mailbox simultaneously.

Addressing PST Files

While personal archives (and large mailboxes) eliminate the need for PST files on a go-forward basis, turning on a user's personal archive does not, by itself, address the problem presented by previously created PST files—and does not prevent users from continuing to use them.

Consequently, in conjunction with deploying personal archives, the additional steps of "locking down" (turning off) and removing existing PST files should be undertaken—along with preventing the creation of new PST files; not an insignificant technical and operational challenge.

Technical Challenges. Technical challenges include: identifying where the files are stored, and, once they are found, how to ingest them; multiple versions;³ and file formats.

Operational Challenges. Operational challenges are typically culture driven. Taking away PST files is likely to be disruptive to users, and can be perceived as an encroachment on professional autonomy.

Notwithstanding these challenges, eliminating PST files is an essential element of leveraging Exchange 2010's records management and eDiscovery features.

File Remediation. To address PST files, organizations should initiate a PST file remediation project with the goal of permanently eliminating PST files.

Project elements should include:

- Instituting an email retention policy by utilizing the records management features of Exchange 2010 (discussed later in this article). This step should include training users how to move email out of PST

³ PST files created with Outlook 2002 (or earlier versions) were limited to two gigabytes, and if the files exceeded this threshold, they could become unstable and corrupt. Outlook versions 2003 and later have a maximum size of 20 gigabytes. For more information about PST files, see http://en.wikipedia.org/wiki/Personal_Storage_Table

files and into a managed location (where retention policies can be applied).

- Deploying a tool to capture, en masse, PST files from across the organization’s IT landscape. Microsoft’s Exchange PST Capture tool identifies, collects, and ingests PST files found in an organization’s network.⁴
- Making PST files “read-only” so users can access items already in their PST files but cannot add new items.
- Informing users that PST files are eventually going to be prohibited both technically and by policy.
- Disabling *all* methods for creating new PST files,⁵ including:
 - Auto archiving (an Outlook function in which items are automatically moved into a PST file at either pre-determined or ad hoc intervals);
 - Outlook data file: an Outlook function that allows users to create PST file containers. While less technically adept users are not likely to be familiar with this function, arguably minimizing its prevalence, it is a method of creating PST files.

A word of caution: changing the way users interact with and use email is a significant undertaking, so

⁴ Tools are also available from Sherpa Software, Symantec, and others. For an overview of Exchange PST Capture see Microsoft TechNet Library, Exchange Server Tools Documentation, Microsoft Exchange PST Capture at [http://technet.microsoft.com/en-us/library/hh781036\(v=exchg.141\).aspx](http://technet.microsoft.com/en-us/library/hh781036(v=exchg.141).aspx) (last visited 12/20/2012)

⁵ For an overview of PST file remediation, see Microsoft TechNet Library, Exchange Server 2010, Messaging Policy and Compliance, Planning for Messaging Records Management at [http://technet.microsoft.com/en-us/library/dd298089\(v=exchg.141\).aspx](http://technet.microsoft.com/en-us/library/dd298089(v=exchg.141).aspx) (last visited 12/20/2012)

strong and committed executive sponsorship should be secured prior to launching a PST file remediation project.

Messaging Records Management (MRM)

Microsoft uses the term “messaging records management” (MRM) to refer to Exchange 2007 and 2010’s records management features and functions. MRM provides several tools which support or enable application of records management policies and procedures to email and other Exchange based data.

Exchange 2007 Approach

MRM was introduced in Exchange 2007, using an approach based on the concept of “managed folders.”⁶ In the managed folders paradigm, individual mailbox folders are assigned retention policies; the retention policies control how long email (and other items) are retained in the folder; and to further refine retention, content settings can be configured to specify retention for certain items (voice mails, contacts, tasks, etc.).

Users are able to move items into a folder they know to have a specific retention policy (such as “keep for 10 years”); items not moved are retained per the policy applied to the folder in which the item originally appears.

The table below is an illustration of a managed folders retention policy.

⁶ For information regarding Exchange 2007’s approach to MRM, see **MCITP: Microsoft® Exchange Server 2007 Messaging Design and Deployment: Study Guide**, By: Rawlinson Rivera, Publisher: John Wiley & Sons, Pub. Date: January 29, 2008, Chapter 16: Planning Exchange Server 2007 Compliance or TechNet at [http://technet.microsoft.com/en-us/library/bb123507\(v=exchg.80\).aspx](http://technet.microsoft.com/en-us/library/bb123507(v=exchg.80).aspx)

| Mailbox Folder | Folder Type ⁷ | Retention Policy |
|----------------|--------------------------|---|
| Inbox | Standard/Default | Delete items in the Inbox that are more than 180 days old |
| Sent Items | Standard/Default | Delete items in the Sent Items folder that are more than 180 days old |
| Deleted Items | Standard/Default | Delete items in the Deleted Items folder that are more than 30 days old |
| 1 Year folder | Custom Folder | Delete items in the 1 year folder that are more than 1 year old |
| 3 Year folder | Custom Folder | Delete items in the 3 year folder that are more than 3 years old |
| 10 Year folder | Custom Folder | Delete items in the 10 year folder that are more than 10 years old |

⁷ If managed folders are in use, two types of folders may appear in a user’s mailbox: default—which all users have; and, custom—which just a sub-set of users may have.

The managed folders approach to MRM is deprecated in Exchange 2010 in favor of a new approach: retention tags (discussed next).

While managed folders can be implemented in Exchange 2010, there are at least two factors which mili-

tate against using them: first, users assigned a managed folder policy cannot have a personal archive; and second, the managed folder approach to MRM is completely removed in Exchange 2013.

Exchange 2010 Approach

The Exchange 2010 approach to MRM employs the concept of retention “tags.” In the retention tag paradigm, users assign a retention period to email and other items by assigning a retention tag selected from a menu of tags. Items not tagged by a user can inherit a default retention period; default retention periods can be assigned to one or more folders in a user’s mailbox.

Key steps to configure retention tags include:

- Creation of tags. There are three tag types: default, personal, and retention policy.⁸

⁸ *Personal Tags* can be applied to individual items or a user’s custom folders (if they have any.) *Retention Policy Tags (RPT)* are applied to the default folders which appear in each user’s mailbox (Inbox, Sent Items, Deleted Items, etc.) *Default Policy Tags (DPT)* are applied to items that are either a) not assigned a Personal Tag or b) are not in a folder that has a Personal Tag assigned. For a detailed explanation, see Microsoft TechNet, Exchange 2010, Messaging Policy and Compliance, Messaging Records Management Strategy at <http://>

- Configuration of tags: what should happen to items based on the tag selected? Tags can be configured to invoke one of the following actions:
 - o Move to archive;
 - o Delete and allow recovery;
 - o Permanently delete; and
 - o Mark past retention limit.
- Assignment of tags to a retention policy: once tags have been created and configured with the appropriate retention and disposition options, they are assigned to a retention policy.
- Retention policies are assigned to the mailboxes.

The table below is an illustration of a retention policy using the retention tag based approach to MRM (adapted from a table found at technet.microsoft.com).⁹

[technet.microsoft.com/en-us/library/dd297955\(v=exchg.141\).aspx](http://technet.microsoft.com/en-us/library/dd297955(v=exchg.141).aspx) (last visited 5/7/2013.)
⁹ See [http://technet.microsoft.com/en-us/library/dd297955\(v=exchg.141\).aspx](http://technet.microsoft.com/en-us/library/dd297955(v=exchg.141).aspx) (last visited 12/18/2012)

| Retention Tag Name | Tag Type | Retention Period in Days | Action |
|---|--------------------------|--------------------------|---------------------------|
| Default 2 years move to archive | Default Policy Tag (DPT) | 730 | Move to Archive |
| Personal 1 year move to archive | Personal tag | 365 | Move to Archive |
| Personal 5 year move to archive | Personal tag | 1,825 | Move to Archive |
| Recoverable Items 14 days move to archive | Recoverable Items folder | 14 | Move to Archive |
| 1 Year Delete | Personal tag | 365 | Delete and Allow Recovery |
| 5 Year Delete | Personal tag | 1,825 | Delete and Allow Recovery |
| Never Delete | Personal tag | Not applicable | Delete and Allow Recovery |

Litigation Hold

2010 is the first version of Exchange to *include* litigation hold functionality. Previously, implementing a litigation hold in Exchange required custom programming or deployment of a third party solution.

Key elements of the litigation hold feature in Exchange 2010 include:

- Through a simple graphical interface, individuals assigned the appropriate rights can place any number of mailboxes on litigation hold.¹⁰
- Holds apply to the entire mailbox (all content); while duration of the litigation hold can be configured,¹¹ no controls exist to limit the hold to just certain items or content—it’s an all (the entire mailbox) or nothing approach.

¹⁰ There are three ways to turn on litigation hold: Exchange Management Console (EMC), Exchange Control Panel (ECP), or Exchange Management Shell (EMS.)

¹¹ Length of the hold has to be set using EMS.

- Once a litigation hold is active, items in a user’s mailbox are retained unaltered until the hold is lifted. This element of the litigation hold feature is enabled by the recoverable items folder,¹² a special purpose mailbox folder hidden from users. The recoverable items folder is a repository that captures items users either delete or modify—thus ensuring all items are available for discovery.¹³

¹² The Recoverable Items folder replaces for what was known as the “dumpster” in pre-2010 versions of Exchange. NOTE: some resources refer to the recoverable items folder as the dumpster or “dumpster 2.0.” Unlike the (old) dumpster, the Recoverable Items folder is an actual (albeit hidden) folder which serves an important role in the eDiscovery features (pre-Exchange 2010 the dumpster was not an actual folder, but rather, just a status conferred on items). For details about the Recoverable Items folder see “Litigation Hold and the Recoverable Items Folder” available at [http://technet.microsoft.com/en-us/library/ee861123\(v=exchg.141\).aspx](http://technet.microsoft.com/en-us/library/ee861123(v=exchg.141).aspx) (last visited 5/6/2013.)

¹³ The Recoverable Items folder has a default quota of 30 gigabytes (may be increased or decreased). Once the quota is

- Turn on or off the optional litigation hold alert feature. If the alert feature is enabled, a message (alert) is displayed in the “backstage”¹⁴ area of Outlook (requires Outlook version 2010 or higher). The alert can include a link to additional resources, such as an intranet page containing details about the litigation hold.
- Auditing: as discussed below, user and administrator actions are logged in Exchange 2010. Turning on and off a litigation hold is one of the actions logged and available for reporting.
- Reporting: reports are available showing who is on litigation hold, when the hold was enabled, and who turned on the hold.

Journaling

Exchange journaling captures a copy of all incoming and outgoing messages for some or all users. The copy is placed in either a separate “journal” mailbox or a third-party archiving tool. Journaling ensures that, regardless of end-user attempts to delete messages, a copy of each message is captured.

Introduced in Exchange 5.5, journaling was enhanced in Exchange 2007 with the addition of “premium” options. Prior to Exchange 2007, journaling could only be performed at the database level, meaning all users in a database had to be journaled; there was no ability to select individual users—potentially leading to over-preservation.

However, with the premium journaling introduced in Exchange 2007 and now available in Exchange 2010, specific users located in any database could be selected for journaling. In addition to being able to target individuals for journaling, premium journaling also allows for fine tuning of what messages get captured: global (all messages coming in and going out regardless of origin); internal (only messages originating within an organization); and external (only messages originating from outside an organization.)

Considerations pertinent to journaling include:

- Who should be journaled? All users or just a subset? Some organizations journal just executives, while other organizations journal all users. In some situations journaling will be applied to users on litigation hold to preserve their email on a go forward basis.
- What should be journaled? All messages? Only messages sent within the organization? Only messages originating or going outside the organization? Should voice mail messages and missed call notifications be journaled?
- Where should journaled email go? In many cases, messages will be copied first to an Exchange journal

reached, items may no longer be added. Since the Recoverable Items folder is a key enabling component of Exchange 2010’s litigation hold function (it is the repository in which deleted and altered items reside), ensuring the quota is not exceeded is important.

¹⁴ The backstage is a navigational element of Outlook 2010 incorporating some of the features and functions formerly available via the file and tool menus. The backstage presents a full screen view of the menu items, features, and functions it makes available. If a user is on litigation hold, a message can (optionally) be displayed in the backstage area. For example, a message with a link to web site can be provided to assist users on litigation hold.

mailbox and then transferred into a third party archiving tool.

- Each journaled message is delivered as an attachment to a “journal report.” The journal report provides information about the message such as the sender and recipient email addresses.
- If Microsoft Information Rights Management (IRM) is used, then Journal Report Decryption should be enabled; if it’s not, IRM protected messages that have been journaled may not be accessible—which could lead to problems in an eDiscovery context.

Discovery Search

For legal and RIM professionals, perhaps the most significant new feature in Exchange 2010 is discovery search. Discovery search enables a user with appropriate rights to conduct keyword and metadata parameter searches across an organization’s mailboxes. From within Exchange 2010 and without any additional software, legal and RIM professionals can search for and retrieve information potentially responsive to any type of investigation, litigation, or regulatory matter.

While discovery search is a major step forward for Exchange in terms of out-of-the-box functionality, as discussed next, it has limitations with which legal and RIM professionals should be familiar. Whether discovery search is a viable legal discovery tool for identifying and retrieving email and other Exchange/Outlook artifacts is a question that can be answered only by legal and RIM professionals familiar with the details of how it works.

Main Elements. The main elements of Exchange 2010 discovery search are:

- **Assignment of the discovery management role.** To conduct discovery searches, a user must be assigned to the Discovery Management role. The purpose of the Discovery Management role is to allow non-IT staff to perform discovery tasks such as conducting discovery searches and enabling litigation holds.
- **Access method.** While discovery search can be accessed via the Exchange Management Shell (EMS) using Windows PowerShell Cmdlets, it is highly unlikely non-technical staff will do so (despite the availability of more search and reporting options via EMS.)¹⁵ Those assigned the Discovery Management role will most likely access discovery search via the Exchange Control Panel (ECP).
- **Mailbox locations to search.** By default four mailbox locations are searched:¹⁶ 1) standard mailbox; 2) personal archive (if one is present); 3) standard mailbox Recoverable Items folder; and 4) personal archive Recoverable Items folder.¹⁷
- **Search query parameters.** Discovery search allows search queries to be built using the following search parameters:

¹⁵ For a discussion about conducting discovery searches using Windows PowerShell, see **Microsoft Exchange 2010 PowerShell Cookbook** By: Mike Pfeiffer; Publisher: Packt Publishing Pub. Date: July 25, 2011 page 343 (Performing a Discovery Search)

¹⁶ When conducting a search via the Exchange Management Shell (EMS), the Recoverable Items folder can be excluded.

¹⁷ The mailbox and personal archive each have a Recoverable Items folder.

- o Object types: by default only email messages (and—see below—some attachments) are searched. Other Exchange and Outlook objects (meetings, tasks, notes, journal, contacts, and (if enabled) instant messaging conversations) can be included in a search.
- o Specific email addresses or domains: search for email messages sent to or from a specific email address (all messages from jdoe@abc.com), or to or from a specific domain (all messages to or from the “microsoft.com” domain).
- o Date range: search for messages falling into a specific date range. For example: all messages after 10/1/2010, or, all messages between 10/1/2009 and 10/31/2012.
- o Keywords: search for messages containing only certain words, phrases, or combinations of words and phrases. The keyword search function in Exchange 2010 is not as robust or sophisticated as that found in third party eDiscovery tools. While Exchange 2010 does allow for keywords to be used with standard Boolean operators AND, OR, and NOT, more sophisticated searches require expertise with Microsoft’s AQS search engine.¹⁸
- o Include unsearchable items: due to limitations in Exchange 2010’s ability to index—and make searchable—certain types of content, discovery search includes a feature whereby non-searchable items may be included or excluded in the search results. Unsearchable items are discussed in more detail below.
- o Limit search by mailbox: searches can be conducted of single, multiple, or all mailboxes.¹⁹
- **Deduplication.**²⁰ If enabled, deduplication filters search results so that only one copy of each responsive item appears in the search results, regardless of how times the message appears in the source mailbox or mailboxes. If full logging is enabled for a search, a report listing each instance of deduplicated messages is delivered with the search results. If deduplication is not enabled, then each instance of the message is delivered, along with the folder and parent folders (if any) in which the message appears.
- **Search results delivery.** Where discovery search results are delivered is determined by how the discovery search is conducted, as follows:
 - o **Exchange Control Panel (ECP):** for discovery searches executed through the ECP, the only option for delivery is to a discovery mailbox (discovery mailboxes don’t belong to any one user, and are accessible only to those serving in the discovery management role).²¹

¹⁸ Advanced Query Syntax (AQS) is the syntax (language) used in the Windows Search and Windows Desktop Search (WDS) technologies. For a description of AQS and examples of the syntax see <http://msdn.microsoft.com/en-us/library/aa965711.aspx> (last visited 4/19/2013).

¹⁹ The maximum number of mailboxes which can be searched concurrently is 25,000. If more than 25,000 mailboxes need to be searched, multiple searches are required. See “Performing a Discovery Search” at [http://technet.microsoft.com/en-us/library/dd335072\(v=exch.141\).aspx](http://technet.microsoft.com/en-us/library/dd335072(v=exch.141).aspx) (last visited 5/5/2013).

²⁰ Deduplication requires Exchange 2010 SP1 or higher.

²¹ A single discovery mailbox is created by default when Exchange 2010 is deployed (additional ones may be created). Discovery mailboxes are hidden accounts, meaning they don’t show up in Exchange address books. Otherwise, they function like regular mailboxes, and don’t have any special capabilities.

o **Exchange Management Shell (EMS):** discovery searches executed via EMS can be delivered to any mailbox; the target does not have to be a discovery mailbox, although it could be.

- **Search log:** every discovery search results set includes a log containing information about the search. The amount of detail in the log depends on whether basic or full logging is selected:

- o Basic logging captures: date and time the search was started and finished, who executed the search, size of the result set, keyword hits (number of hits and number of mailboxes), and mailboxes searched.

- o Full logging captures: includes all basic logging information plus a report (in CSV format) with details about each item retrieved by the search.

- **Export for delivery:** once a search is completed, the next step is to deliver the results to the next stage of the process. Discovery search results can be moved or copied into a PST file (this must be done within Outlook client software; exporting to a PST file is not available from within OWA). As soon as the results have been put into a PST file, the PST file can be transported using regular file transfer methodologies.

Indexing Technology. A critical component of discovery search is the underlying indexing technology (the indexing technology is what renders content searchable). In Exchange 2010, indexing is performed by the built-in Exchange Search indexing technology.

Legal and RIM professionals should be familiar with the indexing limitations of Exchange Search, as the efficacy of identification and preservation executed via discovery search is dependent on the quality of Exchange Search indexing.

Not identifying and preserving email and other content due to weaknesses in the technology could open the window to claims of inadequate search and preservation efforts.

With this in mind, consider the following overview of Exchange Search’s indexing and search²² capabilities:

- **What is always indexed?** Outlook user artifacts.

- o For email messages, the following fields: To, From, cc, Bcc, subject, date, message body, message expiration date.

- o For appointments, events, or meetings, the following fields: start time, end time, location, etc.

- o Contacts, tasks, notes, and journal items also have a number of fields which are indexed.

- **What “may” be indexed?** Attachments are most frequently found in email messages, but can also be found with Outlook artifacts (for example, a meeting agenda attached to a calendar invitation).

- o Attachments: may or may not be indexed depending on whether an iFilter for the attachment’s file type is installed. By default, iFilters are installed for Microsoft Office (Word, Excel, PowerPoint) and other file types (html, rtf, txt, and so on.) What should leap out at legal and RIM professionals is that PDF files are not, by default, indexed. To index PDF files in Exchange 2010 requires installation of a PDF iFilter (one is available at no charge from Adobe).²³

²² In order to be searchable, items must first be indexed.

²³ iFilters for some non-Microsoft applications may be available from third parties. For example, Adobe offers a free

§ It is imperative to determine whether iFilters are available for the file types the organization wants to ensure are indexed. For example, if an organization uses computer-aided design (CAD) software, and desires files created with the software to be indexed by Exchange Search, it should be determined if an iFilter exists for that particular CAD file type.

o Information Rights Management (IRM)²⁴ protected content. IRM is a Microsoft technology which allows access controls to be placed on messages and attachments. In the context of Exchange, IRM can restrict the forwarding, copying, or printing of email messages (and message attachments.) In addition, messages (and attachments) can be set to “expire” so as to not be accessible after a designated point in time. IRM protected messages are indexed unless originating from an Active Directory forest other than the one where the discovery search is being run.²⁵

■ **What is not indexed?**

o S/MIME encrypted messages.²⁶

o “Safe List” file types. Microsoft maintains a so-called “safe list”²⁷ of file types whose content cannot be indexed. Many of the safe list file types are graphics and multimedia files, such as bmp, jpeg, mp3, and so on. File types appearing on the safe list are not likely to have an iFilter available; Exchange classifies file types on the safe list to be unsearchable.

■ **Unsearchable items folder:** discovery searches can be configured to include unsearchable items; if so configured, items not indexed (excluding safe list items) are returned in the search results. This means S/MIME encrypted messages and items that don’t have an iFilter installed (but are not on the safe list) will be returned in the results set. Including unsearchable items in a search ensures items that would not be identified via discovery search are nevertheless included in the results and thus can be reviewed to determine responsiveness.

Auditing

During regular day to day operation, Exchange executes hundreds of actions, processes, and functions.

iFilter for PDF files. See <http://www.ifilter.org> for information about iFilters; www.ifiltershop.com is a web site offering iFilters for sale.

²⁴ For information about Information Rights Management (IRM) in an Exchange environment, see [http://technet.microsoft.com/en-us/library/dd638140\(v=exch.141\).aspx](http://technet.microsoft.com/en-us/library/dd638140(v=exch.141).aspx) (last visited 5/9/2013.)

²⁵ Even a brief explanation of Active Directory concepts such as forests is beyond the scope of this article. However, it should be noted that the larger and more distributed an organization, the more likely their environment may encompass multiple forests—and consequently some or all IRM protected messages may not be indexed. See **Active Directory for Dummies, 2nd Edition** by Steve Clines and Marcia Loughry for a good introduction to Active Directory terminology and concepts.

²⁶ S/MIME is an internet protocol for encrypting (and digitally signing) data. Encryption typically involves use of a “key” to lock and unlock content (S/MIME works on this principle); loss of the key poses risk as the encrypted data is not likely to be retrievable if the key is lost; data that can’t be retrieved because of a lost key might be viewed as data that has been destroyed.

²⁷ The safe list is available at [http://technet.microsoft.com/en-us/library/ee633485\(v=exch.141\).aspx](http://technet.microsoft.com/en-us/library/ee633485(v=exch.141).aspx) (last visited 4/10/2013).

Whether it’s an email being sent, a mailbox being created, or a retention policy being configured, a log of what was done, when, and by whom, can be maintained.

While logging has long been a feature of Exchange, its purpose was primarily to aid system administrators responsible for ensuring system equilibrium (uptime, optimal resource utilization, overall system efficiency, etc.). Now that Exchange includes records management and eDiscovery functionality, the ability to establish an audit trail of who did what and when is essential.

The intent in this article is to alert legal and RIM professionals to the existence of these auditing functions and describe their basic capabilities. There are two points readers should bear in mind regarding Exchange 2010 auditing.

First: the auditing features are limited—don’t expect the same type of robust and sophisticated reporting tools available in dedicated records management and eDiscovery tools.

Second: using the auditing tools in Exchange 2010 requires a degree of technical skill and in some instances advanced technical skills.

Administrator Auditing

Administrator auditing logs actions taken by administrators in Exchange, such as:²⁸

- Adding mailboxes;
- Deleting mailboxes;
- Changing a user’s role (access rights) within the Exchange 2010 environment; and
- Placing mailboxes on litigation hold.

Several key characteristics of administrator auditing are as follows:

- It is enabled by default.²⁹
- While most activities involving changes initiated by administrators are logged and reported, exactly what gets audited and how much detail is captured, can be configured.
- By default, logs are retained for 90 days (can be adjusted to a maximum of 68 years).

Administrator auditing data can be accessed via Exchange Control Panel (ECP) or Exchange Management Shell (EMS.) ECP provides access to a utility that exports a pre-configured report (permission to generate audit reports must have been assigned to the requestor).

Alternatively, custom reports can be created using EMS; while powerful, custom reports require facility with Windows PowerShell. Both ECP and EMS generated reports are delivered in XML format, and require additional processing to be useful.

²⁸ For details regarding administrator auditing in general, and what gets logged specifically, see [http://technet.microsoft.com/en-us/library/dd335052\(v=exch.141\).aspx#Agent](http://technet.microsoft.com/en-us/library/dd335052(v=exch.141).aspx#Agent) (last visited 4/30/2013)

²⁹ Administrator Auditing is enabled by default in Exchange 2010 SP1 and higher; it is not enabled by default in Exchange 2010 RTM (release to manufacturing).

Mailbox Auditing

Mailbox auditing logs mailbox activities: who did what and when with or to a mailbox. Mailbox actions audited include:³⁰

- User deletion of email;
- Discovery searches; and
- Delegates.³¹ sending email on behalf of other users

Several key characteristics of mailbox auditing are as follows:

- It is not enabled by default;
- It is enabled on a mailbox by mailbox basis;
- Because the volume of actions taken in a mailbox can be significant, defining what to audit is essential. For example: should deleting or moving messages be audited? Should delegates sending messages on behalf of other users be audited?
- What types of users³² should have their access audited? There are three user types for the purposes of mailbox auditing: (1) the owner of the mailbox (traditionally considered the “user”); (2) delegate; (3) administrative
- Mailbox audit logs are retained by default for 90 days, but this can be changed to a maximum of 68 years.

The reporting options for mailbox auditing mirror those available for administrator auditing: generate a pre-configured report from ECP or create custom reports using EMS. As with administrator auditing, the output format is XML.

Transport Rules

Exchange 2007 introduced a new system architecture, the hallmark of which is the employment of five “server roles.”³³ One of these server roles, called the hub transport, operates as a funnel through which all incoming and outgoing messages must pass—which enables Exchange 2010, out-of-the-box, to apply sophisticated processing and filtering rules to those messages.

Transport rules work in two stages: if a condition exists (stage one), an action is taken (stage two).

Transport rules are invoked when messages meet certain conditions. For example: any message sent from a designated user (stage one/condition) is automatically carbon copied (cc) to the user’s manager (stage two/

action); any message *with a pattern of numbers* typically associated with a social security number (stage one/condition) are sent to a designated user for review prior to transmission (stage two/action).

There are hundreds of different conditions and actions which can be used to accomplish an organization’s message filtering objectives.

Elements of Transport Rules. The key elements of transport rules are:

- *Conditions* are attributes or characteristics which, if they appear in a message, trigger the second stage of transport rules (actions) to perform some type of operation on the message. There are some 40 conditions, including:
 - o If a message is from a certain person, persons, or to a particular distribution list;
 - o If a message subject field contains a certain word, words, phrase, number, or matching “expression” (expressions are discussed below);
 - o If an attachment to a message contains a certain word, words, phrase, number, or matching “expression” (expressions are discussed below).
- *Actions* are the operations performed on messages meeting the (stage one) conditions. Sixteen different actions can be invoked, including:
 - o Add text to a message (such as a disclaimer);
 - o Copy (cc or Bcc) the message to certain email address (or addresses);
 - o Forward the message for moderation.³⁴
- *Exceptions.* Further refinement of transport rules is accomplished by employing exceptions which, if present, preclude an action from taking place. Forty different exceptions are available, including:
 - o Except when a message is from a certain person;
 - o Except when a message is marked with particular message classification;
 - o Except when a message contains certain words.
- *Regular expressions* are “a concise and flexible notation for finding patterns of text in a message.”³⁵ Regular expressions parse the character strings in specified content, looking for pattern matches. For example: `\\-\\-\\-\\` is a regular expression that identifies social security numbers (\\ represents any number value between 0-9.) Regular expressions are useful for identifying high-risk content types, such as credit card and social security numbers.³⁶

Scenarios where transport rules could be used include:

- *Disclaimers:* text added to the body of a message which states a policy or qualification regarding the message. Some organizations will append a dis-

³⁰ For a list of actions that can be logged see [http://technet.microsoft.com/en-us/library/ff459237\(v=exch.141\).aspx](http://technet.microsoft.com/en-us/library/ff459237(v=exch.141).aspx) (last visited 4/30/2013)

³¹ A delegate is an Exchange user who has been granted access and certain rights to another user’s mailbox, and can engage in certain mailbox activities such as responding to messages and meeting requests.

³² User types are as follows: 1) mailbox owner is the individual assigned to the mailbox; 2) delegates are users who have their own mailbox (hence are mailbox owners of their own mailbox) but also have access and rights to another’s mailbox; 3) administrators are a third type of user, but their usage differs in that they are not regularly or actively accessing other users’ mailboxes—rather, administrator access is based on some type of administrative function or requirement.

³³ The five server roles are: mailbox, client, unified messaging, hub transport, and edge transport. For an overview of these roles see [http://technet.microsoft.com/en-us/library/dd298026\(v=exch.141\).aspx](http://technet.microsoft.com/en-us/library/dd298026(v=exch.141).aspx) (last visited 5/10/2013).

³⁴ Moderation is an Exchange 2010 feature where messages are submitted to a designated individual for review prior to delivery. If approved, messages are delivered; if not approved, the rejected messages are returned to the sender undelivered. For a detailed discussion of moderation see Microsoft TechNet at [http://technet.microsoft.com/en-us/library/dd297936\(v=exch.141\).aspx](http://technet.microsoft.com/en-us/library/dd297936(v=exch.141).aspx) (last visited 4/16/2013).

³⁵ Microsoft TechNet, Exchange Server 2010, Library at [http://technet.microsoft.com/en-us/library/aa997187\(v=exch.141\).aspx](http://technet.microsoft.com/en-us/library/aa997187(v=exch.141).aspx) (last visited 5/3/2013).

³⁶ For an explanation of expressions and regular expressions in Exchange 2010, see [http://technet.microsoft.com/en-us/library/aa997187\(v=exch.141\).aspx](http://technet.microsoft.com/en-us/library/aa997187(v=exch.141).aspx) (last visited 5/2/2013).

claimer similar to the following on all outgoing messages: “This email message is intended to be received only by persons entitled to receive the confidential information . . .” Application of a disclaimer can be universal (applied to all messages) or tailored (only applied to messages sent outside the organization).

- Ethical walls: certain groups of employees can be configured so they are prohibited from engaging in email communications with other groups of employees.
- Message screening: messages are reviewed and must be approved by designated individuals prior to delivery. For example, any message containing the words *merger* or *acquisition* are forwarded to a manager for approval; the manager then approves or rejects the message. The sender is notified whether the message approved or rejected. (NOTE: this scenario is also referred to as moderation).
- Compliance monitoring: specified users’ messages are cc’d to a designated mailbox for review. Unlike screening where messages are approved or rejected, in a compliance monitoring scenario, messages are delivered as usual, but can be reviewed by compliance staff on a periodic basis to ensure conformance with organizational policies.

Transport rules are a potentially useful tool for enabling a number of email management policies. As with message classifications and auditing, collaboration among IT, legal, and RIM professionals is essential to leveraging transport rules fully.

Message Classifications

Introduced in Exchange 2007, message classification allows for electronic labels (which are then prominently displayed in the message header) to be applied to select messages by users or transport rules (discussed below).

For example, a policy could be instituted by an organization that requires users to apply a message classification of “COMPANY CONFIDENTIAL” to any messages containing proprietary company data.

Users apply message classifications by selecting a label from a menu of choices available in the Outlook or OWA client. Transport rules apply message classifications based on the existence of pre-determined conditions, such as certain words appearing in the subject line or body of a message.

Setting up message classifications includes the following steps:

- Determination of what labels and descriptions suit the organization’s objectives (popular message classifications include “attorney-client privileged,” “company confidential,” “secret,” etc.);
- Labels can only be created using the Exchange Management Shell (EMS);
- Labels are “pushed out” to Outlook users (if users are to be given the ability to label messages) or transport rules are created to stamp messages based on pre-determined conditions.

Since configuring message classifications is done primarily through the EMS, legal and RIM professionals will need to work closely with the appropriate technical resources to leverage this feature.

Message Tracking

While not an eDiscovery or records management feature per se, message tracking (available since the first version of Exchange 4.0) can be a useful tool for tracking a few key pieces of information about messages, including: date and time a message was sent; who sent or received the message; the subject line of the message; whether a message was successfully delivered, and so on.

Message tracking is enabled by default, and has several standard settings which may be of interest to legal and RIM professionals, including:

- Message tracking logs are retained for 30 days (can be changed to a longer time period);
- Message tracking log files are limited to 10 megabytes;
 - o Message tracking log directories are limited to 250 megabytes (these limits may be expanded to the suit the organization’s needs);
- Subject line logging (tracking) is enabled by default, but can be turned off.

There are three ways to access message tracking (all require appropriate permissions).

- Exchange Control Panel (ECP): under “Mail Control,” one of the options is “Delivery Reports.” Clicking on delivery reports brings up a search screen with the following fields available to construct a search: mailbox; sender; recipient; words appearing in the subject line. The results can be viewed on screen, printed, or e-mailed. This is the only access method likely to be used by legal and RIM professionals.
- Tracking Log Explorer: while this option includes more search parameters than ECP, it is an administrator-only tool (it can only be accessed via the Exchange Management Console). Results are delivered in XML format and can be viewed on screen, e-mailed, or printed.
- Exchange Management Shell (EMS): while offering the greatest reporting power and flexibility, EMS requires skill with Windows PowerShell. Legal and RIM professionals will need to work closely with technical staff to leverage this option.

Message tracking is a useful tool for quick and basic investigation of whether a message was sent, when it was sent, to whom, and what words appear in the subject line.

Glossary

Advanced Query Syntax (AQS) AQS is a protocol (syntax) for creating search queries in certain Microsoft applications or systems, including Windows Desktop Search, Windows Search, and Exchange 2010. AQS supports Boolean operators (AND, OR, NOT), wildcards, and metadata (parameter) searches (to, from, subject, etc.) An AQS query might look as follows: from:jcollins@theingersollfirm AND subject:esi data map. This AQS query would find all messages where the from field contains jcollins@theingersollfirm and the subject field contains ESI data map; both conditions must be met because of the AND operator.

Cmdlet Cmdlets are instructions entered into the Windows PowerShell user interface. Cmdlets instruct PowerShell what actions to perform, using a standard

verb-noun format (for example: search-mailbox; search is the action (verb) and mailbox is the what (noun.) Cmdlets can range from the simple (just a few words) to the very complex (a paragraph or more.)

Delegate A feature wherein a mailbox owner grants another user access to his or her mailbox. For example, a company's CEO might grant delegate access to an assistant; the assistant (depending on the level of delegate access granted) could respond to email messages and meeting requests on behalf of the CEO, relieving the CEO of having to respond. Depending on the exact configuration, messages sent to the mailbox owner may or may not be delivered into the owner's mailbox; an option is available where only the delegate receives certain messages (such as meeting requests.) An excellent summary of the delegate feature is available in **Microsoft Outlook 2010 Inside Out** by Jim Boyce, O'Reilly Media, Inc. on behalf of Microsoft Press, 2010, Chapter 34.

Exchange Control Panel (ECP) ECP is a web-based administration console accessed via Outlook Web App (OWA.) ECP provides administrators (and others) with a graphical user interface to configure and manage a number of different features and functions on an organization-wide basis, such as conducting discovery searches, adding, deleting, and managing distribution groups, putting users on litigation hold, configuring mailbox journaling, and so on. Access to ECP in general, and what it controls specifically, is predicated on role assignment (referred to as Role Based Access Control or RBAC.) The ECP can also be made available to non-administrative personnel, but with a significantly different set of features and functions—and the scope of control is only over the user's own environment.

Exchange Management Console (EMC) EMC is one of the two primary administrative interfaces available for Exchange 2010 (the other being EMS). EMC is a traditional, GUI-based client/server application. Most day-to-day and ad-hoc administration tasks can be completed via EMC, including: configuring mailbox storage quotas; configuring retention tags and policies; creating new mailbox databases; adding users, and so on. EMC may be thought of as the interface through which much of the administrative "heavy lifting" is conducted. Legal and RIM professionals are not typically granted access to the EMC.

Exchange Management Shell (EMS) A command-line (think MS-DOS) administrative control tool for Exchange. EMS includes all the functionality found in the Exchange Management Console (EMC) and Exchange Control Panel (ECP) (in fact, any function executed by EMC or ECP is done via EMS running behind the scenes) plus additional functions not found in those two

interfaces. EMS facilitates automation of tasks via scripting (the language used in EMS is Windows PowerShell.) EMS is likely to be used only by system administrators. However, EMS is the most powerful (and complex) management tool available for Exchange 2010.

Information Rights Management (IRM) IRM is a Microsoft technology which enables content created in some Microsoft applications (Outlook, Word, Excel, PowerPoint) to have restrictions and controls embedded; these restrictions and controls determine what recipients can and cannot do with the content. For example, IRM can control whether a user can view, print, copy, or forward an Outlook email message (this control can be extended to any attachments.) Also, content can be assigned an expiration date, after which the content is no longer accessible. IRM protected content is encrypted, so only those granted appropriate rights can access, view, and interact with the content. IRM works in conjunction with Active Directory Rights Management Services (AD RMS.)

Personal archive (also known as archive mailbox) Introduced in Exchange 2010, the personal archive provides users with a server-based storage location for older messages, or messages to which the user does not need instant access. Conceptually, the personal archive may be viewed as a replacement for .PST files. A key difference between personal archive and PST files is that personal archives, because they are stored on the Exchange Server, can be managed by Exchange records management and eDiscovery features, whereas PST files, because they exist outside the Exchange Server, cannot be managed by these features.

S/MIME(Secure/Multipurpose Internet Mail Extensions) A standard internet protocol for encrypting and digitally signing data that is supported by Microsoft Exchange and Outlook. S/MIME encrypted data can only be read by a user who has the appropriate credentials/rights to decrypt the data. Use of S/MIME in an Exchange and Outlook environment requires a certain degree of configuration and set up, especially the digital certificates required to make the process work.

Windows PowerShell A Microsoft system administration tool that employs a command-line scripting language. This command line scripting language employs the concept of cmdlets to instruct applications (such as Exchange 2010) or operating systems (such as Windows Server 2012) what to do. There are hundreds of different cmdlets, and hundreds of different parameters which can be used with cmdlets. A major benefit of PowerShell is the ability to automate (script) routine tasks. Mastering PowerShell is not for the faint of heart.