

Microsoft Exchange 2010 and Archiving: The Complementary Role of Third Party Solutions

A White Paper by Ferris Research

November 2009. Report 826

Ferris Research, Inc.

One San Antonio Place

San Francisco, Calif. 94133, USA

Phone: +1 (650) 452-6215

Fax: +1 (408) 228-8067

www.ferris.com

Table of Contents

Executive Summary	3
Exchange 2010 Archiving: Tutorial	4
Main Benefits	4
User Model	4
Archiving Policies	4
Retention Features	5
Policy Scope and Access Controls	6
Legal Holds	6
E-Discovery Search	6
PowerShell.....	6
Archiving Roadmap.....	7
Microsoft Validates the Archiving Market.....	7
The Complementary Role of Third Party Archiving	8
Compliance.....	8
More Complex Policies	8
Policy Enforcement	8
Types of Archived Material.....	8
Comprehensiveness of Retention	9
E-Discovery	9
Scope of Searches.....	9
Reliability of Search Results	9
Lack of Case Management Tools	9
Greater Granularity for Legal Holds	9
Comprehensiveness of Search.....	10
Miscellaneous	10
Backup & Storage Management.....	10
Cost.....	10
Offline Support.....	11

Executive Summary

This report explains the archiving features of Exchange 2010, and assesses their benefits and the role of third party archiving vendors. The main conclusions are:

- Exchange 2010 will substantially improve the management of PST files.
- Exchange retention policy management is substantially improved, and basic e-discovery services have been added. These enhancements provide rudimentary facilities for compliance and e-discovery purposes.
- Exchange 2010, consistent with its enhancements over the last five years, has substantial performance improvements. Nevertheless, it is unclear whether backup and restore times will be acceptable for large mailboxes.
- Microsoft recognizes that its archiving offering will not satisfy everyone's needs. The company wants to encourage its large partner ecosystem to provide complementary solutions, preferably building on top of Exchange's archiving, retention, and e-discovery infrastructure.
- Third party archiving vendors will continue to enhance Exchange for the foreseeable future, especially in the areas of regulations compliance and e-discovery. They may also serve a valuable role in reducing backup and restore times, by offloading content to external storage.

Exchange 2010 Archiving: Tutorial

On April 19, 2009, Microsoft announced that Exchange 2010, due for release in mid-2010, will have archiving capabilities.

Main Benefits

The main benefits of Exchange 2010 archiving are:

- *Less need for PST files.* For backup-and-restore reasons, and because of the cost of SAN storage, users often maintain a large part of their message store themselves, in local PST files. These PST files are thus out of IT control, and present major backup, compliance and e-discovery problems. Exchange 2010 brings them back under IT control. Users can insert PST content into their archive mailboxes, located on Exchange servers, thus removing the need for local PSTs.
- *Improved retention.* Exchange 2007 had rudimentary retention policy support. With Exchange 2010, applying retention policies is much easier and more natural.
- *Basic e-discovery services.* A legal hold can be applied to a user's mailbox, and litigation support staff can conduct searches across multiple users' mailboxes.

User Model

The main new feature of Exchange 2010 is that each user can now have a secondary mailbox that contains the archive for that user. This is accessed through Outlook 2010 or the latest version of Exchange's Web client, Outlook Web Access. In Exchange 2010, the Web browser has been renamed Outlook Web App for branding reasons.

Browsing the archive is just like browsing a normal Outlook mailbox. You navigate a hierarchy of nested folders.

The archiving works with any Exchange content, including email, tasks, address book entries, calendar meetings, and notes. Office Communications Server instant messages are also supported. Unlike many third-party archiving systems, it is not just an email archive, although clearly for most people the email archive will be the most important element.

It's striking that archiving is a seamless extension of Exchange and Outlook. As you would expect with a built-in capability, few new concepts are introduced. Users and administrators rely on familiar interfaces, minimizing the learning curve for deploying the new services.

Archiving Policies

Content is inserted into the archive in two ways.

Users can simply drag content (typically email messages or folders containing other folders or email messages) from their primary mailboxes or PST files.

Alternatively, IT can define rules that allow Exchange content to be moved automatically into the user's archive and out of the user's main store. The rules are very simple and are based on time. For example:

- Move to Archive after 30 days.
- Move to Archive after 90 days.
- Move to Archive after 365 days.

The rules all have the format *Move to Archive after <specified time period>*. IT can provide the appropriate retention policies for its users. Users in turn can choose whether to apply the retention policies that IT makes available. The archiving policies kick in automatically, moving content after the appropriate time to the archive mailbox and deleting it from the primary mailbox.

IT also now has better control over deleted email. Past versions of Exchange have allowed savvy users to delete a message, then empty both their deleted items and the deleted item retention (dumpster) location to hide a message. Now IT can force all email to pass through a new and improved "dumpster 2.0" so that intentionally purged email can no longer fly under the radar.

Retention Features

Exchange 2007 introduced rudimentary retention policy support. Administrators can define several folders, such as "Keep for 30 days" and "Keep for 7 years." After the appropriate time, the content in the folders is then deleted.

Users have to abide by the folders defined by IT, and many find this inconvenient. Exchange 2010 has a much more flexible and natural approach. IT can define simple retention rules that automatically delete content after a specified period of time, such as:

- Keep for 90 days
- Keep for 2 years
- Keep for 5 years
- Keep for 7 years

The rules all have the format *Keep for <specified time period>*. Retention policies can be applied to any items, whether they are in the primary mailbox or the archive mailbox. When a message with its own retention policy is inserted into a folder that already has a defined retention policy, the longer retention period applies.

Policy Scope and Access Controls

Move to Archive and *Keep for* rules can only be defined by administrators. Default policies can be applied to individual items or folders.

IT also has the option to define user default policies. For example, users might be given a default policy for their Inbox requiring that content be deleted after six months. Conversely, users have the option of overriding a system default with other IT-defined policies.

Move to Archive and *Keep for* rules can be applied to individual items or to folders that contain content or other folders.

Overall, IT has a lot of flexibility in determining the mix of *Move to Archive* and *Keep for* rules that are available to users, and can apply the rules with a high level of granularity.

Legal Holds

IT can apply legal holds to entire mailboxes so content cannot be destroyed. Users may or may not be notified of such holds, depending on how IT configures the holds. The hold feature lacks the granularity of the *Move to Archive* and *Keep for* policies, applying only at the level of entire mailboxes.

In general, legal holds are for an indefinite period, although it's possible to define the length of a hold.

E-Discovery Search

A Web-based search tool allows searches spanning multiple mailboxes. Search criteria are powerful. For example, you can build up searches from *ands*, *ors*, and *nots*; search works across all Exchange content types and attachments; search OCS instant messages; and search voice messages using the voice-to-text feature. Obviously, search applies to ordinary material as well as material subject to legal hold.

For audit purposes, a searchable log is kept of all e-discovery searches.

This is a substantial enhancement for Exchange. Previously, such searches were very rudimentary and required a technical person who was familiar with the Exchange environment and the use of the ExMerge utility and PowerShell. Now searches are much easier and can be performed by litigation support staff.

PowerShell

In Exchange 2010, you often have to issue commands using the PowerShell command language. This will be rectified over the next 12 months, the normal administrative GUIs will work as well as PowerShell.

Archiving Roadmap

Specific additional enhancements are under way. Microsoft discussed these with Ferris Research, mostly on a nondisclosure basis.

Microsoft recognizes that its archiving offering will not satisfy everyone's needs. For example, many organizations have demanding compliance and e-discovery needs. Microsoft wants to encourage its large partner ecosystem to provide corresponding solutions that complement the native Exchange facilities, preferably building such solutions on top of Exchange's archiving, retention, and e-discovery infrastructure.

Consequently, Exchange 2010 now includes a web service API for e-discovery searches, and other web service APIs are under development.

Microsoft Validates the Archiving Market

Previously, archiving technology has been sought mainly by large organizations, legal firms, and medium-size firms in special markets (for example, healthcare).

By entering the market, Microsoft has validated it. This will educate the market on the needs for archiving, and its benefits. The general affect will be to encourage archiving adoption, and the use of third party archiving tools that complement Microsoft's own offering.

The Complementary Role of Third Party Archiving

Exchange 2010 archiving will not replace third party archiving tools. It's more accurate to see third party archiving solutions as Exchange enhancements. They will continue to have a synergistic relationship with Exchange for the foreseeable future, rather than a competitive one.

Here we discuss the main ways in which third party archiving tools can enhance Exchange 2010's built-in archiving.

Compliance

More Complex Policies

The structure and meanings of rules for *Move To Archive*, *Keep for*, and legal holds are rudimentary and will be insufficient for many compliance policies.

For example, users who do little foldering probably have very large Inbox and Sent folders, and might want to deploy the following rules:

- Move to Archive if Subject line contains "Archive", or if To/From address contains a competitor's domain name.
- Keep for 7 years if email To or From address is in the Finance distribution list, and if the body or attachment contains "Annual Report" or "Annual Return."

Policy Enforcement

IT defines the available retention policies and can define default *Move to Archive* and *Keep for* policies. In principle, IT can impose policy by defining default archive and retention policies, and then not providing any additional policies.

However, in practice users will often be able to apply alternative policies, and thus they will have the power to decide what to archive and how long to keep such material.

For many regulations, this is inappropriate. Third party solutions can help ensure that policies can be formulated by central compliance staff and be automatically enforced, without giving users the ability to disobey the policy. When they search an archive, auditors and investigators need to be able to have confidence in their results.

Types of Archived Material

The archiving covers all Exchange content types, and Office Communications Server/OCS instant messages.

Some third party archiving solutions cover a broader range of electronic content. The ability to archive ordinary files and SharePoint content are especially valuable.

Comprehensiveness of Retention

Users can manually ingest PSTs into the archive. However, as a practical matter, they often won't know how to do this, or will miss a PST, or simply won't perform the ingestion. Thus it's hard to be sure that retention policies are being applied to all relevant material. There are likely to be PSTs that have been overlooked.

What's needed are tools which automatically discover PSTs (ie, PST crawlers), and automatically ingest them into the archive.

E-Discovery

Scope of Searches

We've noted that Exchange searches are limited to Exchange content types, and OCS instant messages. E-discovery searches should preferably encompass other types of electronic information, including flat files and SharePoint content.

Reliability of Search Results

An Exchange 2010 mailbox is under user control until a legal hold is place. In the interim, users can edit content, or delete it.

Given that archiving is not automatic, the material put on hold may therefore not reflect the original content. Third party tools can help, by enforcing the archiving of material.

Lack of Case Management Tools

Tools to support e-discovery searches are limited.

For example, search results are normally exported to a mailbox. This creates additional copies of information, which is then itself liable to e-discovery. Sifting through the mailbox to whittle down the search set is time-consuming compared to other interfaces, and such whittling-down is not audited (although the initial search is audited).

Greater Granularity for Legal Holds

As we noted, legal holds operate at the level of entire mailboxes. Many organizations need greater granularity and require holds to be definable across multiple mailboxes.

For example, they might want to apply a hold to all material associated with such-and-such senders and/or such-and-such recipients and/or anyone in the sales department between June 2007 and February 2008.

Putting more on legal hold than is required is not a good idea. It leads to further scrutiny, the delivery of more information than necessary to the other side, and to expanded requests from the other side.

In addition:

- Applying holds at the level of mailboxes means that for many organizations, the mailboxes of senior executives (and others in relevant litigious roles) have the potential to be under constant hold. In effect, no email for such people will ever be expunged. In the worst cases, almost all of an organization's email could be held for litigation, defeating the purpose of disposition policy.
- It can become very difficult to maintain holds if there are multiple holds that overlap different time periods. Some executives in litigious businesses may end up having their mailboxes on permanent hold.

Comprehensiveness of Search

As we noted above, the lack of automatic PST ingestion tools means that when conducting e-discovery, some PSTs may have been overlooked.

Miscellaneous

Backup & Storage Management

A major motivation to adopt email archiving has been to remove content from Exchange databases, in order to accommodate backup windows and recovery times.

Users' primary mailboxes will continue to grow. And now, Exchange databases are set to increase very substantially, as they ingest PST files. On top of this, single instance storage has been removed from Exchange.

Microsoft has been working hard over the last five years to increase the practical size of user mailboxes. It believes that with Exchange 2010 it is now practical for users to have primary mailboxes of 10GB or so, and archive mailboxes of similar size. This jump is due to the mailbox resiliency features of Exchange 2010, and the associated I/O improvements and Database Available Group/DAG infrastructure.

We don't doubt that mailboxes can now be much larger. Our concern is that the increase in email traffic is outflanking the architectural gains. This is a common problem for archiving vendors.

Thus we would not be surprised if message store size remains an issue for Exchange 2010. If so, the ability of third party products to offload content to an external store will be of ongoing value.

Cost

Users can access Exchange 2010 archiving via Outlook Web App and their Web browsers. However, most organizations will want to have Outlook 2010 on the desktop to use Exchange 2010's archiving services. In short, many customers will have to wait until they have an expensive, general system refresh.

Third party tools can provide archiving services in the interim. Some are very inexpensive.

Offline Support

Exchange provides offline support for the user's main mailbox, through OST files. However, no local copy is maintained of the archive mailbox. So users have no access to the archive while offline.

It's not clear how much of a problem this will be. Offline access to an archive is important if you have an unduly small mailbox. But if, as Exchange 2010 promises, you really can have a large primary mailbox, the need for offline archive access may be much diminished.

Author: David Ferris

Editor: Mona Cohen

Sponsorship of This White Paper

Metalogix Software commissioned this white paper with full distribution rights. You may copy or freely reproduce this document, provided you disclose authorship and sponsorship and include this notice. Ferris Research independently conducted all research for this document and retained full editorial control.

The logo for Metalogix, featuring the word "metalogix" in a lowercase, sans-serif font. The "meta" part is in black, and the "logix" part is in orange.

Metalogix Software
890 Winter Street, Suite 115
Waltham, MA 02451
Tel: 877-450-8667

Ferris Research

Messaging. Collaboration. Compliance. Ferris Research analysts bring more experience in these areas than any other firm. Period.

Major areas of interest are email, archiving, e-discovery, information leak prevention, unified communications, instant messaging, SharePoint, and mobile communications. We help:

- IT staff evaluate, implement, and maintain these technologies
- Vendors understand the marketplace and its technologies; explain their products or services to the marketplace; and find strategic partners, raise funds, or sell their company
- Investors find and evaluate investment opportunities

We've been in business since 1990—longer than any other analyst firm in our field:

- Clients include many of the world's largest organizations as well as computer vendors from major corporations to small startups.
- We have published more than 200 [formal reports](#) and 1,100 [short bulletins](#).
- Our [news service](#) has approximately 10,000 readers and covers more than 2,000 highly specialized announcements annually.
- Our [research team](#) shares many decades of experience in our core competencies.

In short, our technology and industry depth helps you understand today's products, where they've come from, where they're going, and their value.

Ferris Research is located at One San Antonio Place, San Francisco, Calif. 94133, USA. For more information, visit www.ferris.com or call +1 (650) 452-6215.

Free News Service

Ferris Research publishes a free daily news service to help you keep current on messaging, collaboration, compliance, and related topics. To register, go to www.ferris.com/forms/newsletter_signup.php. In addition to our daily electronic newsletter, you will receive periodic emails announcing new Ferris reports or Webcasts. To opt out and suppress further email from Ferris Research, click on the opt-out button at the end of each email.