Ferris Research

Microsoft Exchange 2010 Archiving and the Value of Third-Party Solutions

A White Paper by Ferris Research

November 2009. Report 826

Sponsored by:

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

Additional sponsors include:

Table of Contents

Executive Summary

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

- Exchange 2010 will substantially improve the management of personal storage table (PST) files.
- Exchange retention policy management is substantially improved in Exchange 2010, and basic e-discovery services have been added. These enhancements provide rudimentary facilities for compliance and e-discovery.
- Consistent with its enhancements over the last five years, Exchange 2010 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 15, 2009, Microsoft announced that Exchange 2010 will have archiving capabilities. See the press release at www.microsoft.com/presspass/press/2009/apr09/04-15Exchange2010PR.mspx. These capabilities are explained below.

Main Benefits

The main benefits of Exchange 2010 archiving are:

- *Less need for PST files*. Because of backup-and-restore issues and the cost of storage area networks (SANs), users often maintain much of their message stores themselves in local PST files. These PST files are out of IT's control and present major backup, compliance, and e-discovery problems. Exchange 2010 brings them back under IT's 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 the ability for each user to have a secondary mailbox containing the user's archive. This archive is accessed through Outlook 2010 or the latest version of Exchange's Web client. For branding reasons, the Web browser client in Exchange 2010 has been renamed Outlook Web App.

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

The archiving works with any Exchange content, including email, tasks, contacts, calendar meetings, and notes. Office Communications Server (OCS) 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 simple and 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 archiving policies for its users. Users in turn can choose whether to apply the archiving policies that IT makes available. The archiving policies kick in automatically, moving content after the appropriate time to the archive mailbox.

Retention Features

Exchange 2007 introduced rudimentary retention policy support, known variously as message records management or managed folders. Here, 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 must abide by the folders that IT defined, 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.

The retention policies clearly lack granularity. Nevertheless, they may suffice for organizations that have not yet defined their retention policies because of the difficulty in cross-departmental decision making. In the interim, many organizations deploy an all-encompassing policy, such as *Keep for 7 years* or whatever is the maximum retention time required in their industry.

Policy Scope and Access Controls

In Exchange 2010, *only administrators can define Move to Archive* and *Keep for* rules. Default policies can be applied to individual items or folders.

IT can also 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 can override 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. Rules are applied flexibly from the mailbox level, through arbitrary folders, and down to the level of individual items.

Legal Holds

With Exchange 2010, IT can apply legal holds to entire mailboxes so content cannot be destroyed. Any deleted or edited items are retained. A hold may or may not be visible to users, depending on how IT configures it. 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 is possible to define the length of a hold.

Past versions of Exchange have allowed savvy users to hide a message by deleting it and then emptying both their deleted items and the recoverable items store (dumpster). The new legal hold facility prevents users from getting rid of email in this way.

E-Discovery Search

A Web-based search tool in Exchange 2010 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.

Search privileges can be delegated to appropriate audit, compliance, and legal support staff using the new role-based access control (RBAC).

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

This is a substantial enhancement for Exchange. Previously, such searches were rudimentary and required a technical person who was familiar with the Exchange environment and the use of the ExMerge utility and PowerShell. These searches were highly disruptive for IT staff.

PowerShell

For certain purposes, the PowerShell command language is an attractive means of accessing Exchange services.

Nevertheless, the normal GUI approach is often preferable. About 20% of the archiving, retention, and e-discovery features are currently only accessible using PowerShell. Over the next 12 months, these services should also be accessible via GUIs.

Synergy with Third-Party Archiving Vendors

Microsoft's goal is not to take over the archiving world. It will target Exchange users who do not already have archiving, which is about 80% of the total. Microsoft also recognizes that its archiving offering will not satisfy everyone's requirements. For example, many organizations have demanding compliance and e-discovery needs.

The company therefore wants to encourage its large partner ecosystem to provide archiving 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 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 was sought mainly by large organizations, legal firms, and medium-size firms in special markets (for example, health care).

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

The Value of Third-Party Archiving

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

Here we discuss the main ways 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."

Policies Can Be Enforced

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 polices, and then not providing any additional policies.

In practice, however, users will often be able to apply alternative policies and thus 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 automatically enforced without giving users the ability to disobey the policy. When auditors and investigators search an archive, they must be able to have confidence in their results.

More Types of ESI

The archiving covers all Exchange content types and OCS instant messages.

Some third-party archiving solutions cover a broader range of electronically stored information. The ability to archive ordinary files and SharePoint content is especially valuable.

Ingest All PSTs

Users can manually ingest PSTs into the archive. However, as a practical matter, they often will not know how to do this, will miss a PST, or simply will not perform the ingestion. Thus it is hard to ensure that retention policies are being applied to all relevant material. Some PSTs are likely to be overlooked.

Many third-party products have tools (for example, PST crawlers) that automatically discover PSTs and ingest them into the archive.

E-Discovery

Search Encompasses More Types of ESI

We have 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.

Greater Completeness of Search Results

An Exchange 2010 mailbox is under user control until a legal hold is placed on it. In the interim, users can edit or delete its contents.

Material existing after a legal hold is applied will, of course, show up in searches. However, given that preservation is not automatic, material that exists before the hold is applied may have been deleted or changed and thus may not be exposed by a search.

Third-party tools can help ensure the completeness of e-discovery by enforcing the archiving of material.

Improved Case Management

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 narrow the search set is time-consuming compared with other interfaces. And unlike the initial search, it is not 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 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) could 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.

Ingest All PSTs

As mentioned, the lack of automatic PST ingestion tools means that when conducting e-discovery, some PSTs have probably been overlooked. They can come back to haunt you later.

Miscellaneous

Backup and Storage Management

A major motivation to adopt email archiving has been to remove content from Exchange databases 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 Exchange 2010 makes it 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 do not 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.

Reduced 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 isn't 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

Waterford Technologies 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.

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.