


# **Best Practices**

## *for Message-level Exchange Recovery*

---

White Paper

*Backup &  
Recovery*



This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. However, because of the possibility of human or mechanical errors, Aelita Software does not guarantee the accuracy, adequacy, or completeness of any information in this publication, and is not responsible for any errors or omissions or the results obtained from use of such information.

Unless otherwise noted, the example companies, organizations, products, people, and events depicted herein are fictitious. No association with any real company, organization, product, person, or event is intended or should be inferred.

Aelita Software does not endorse or accept any responsibility for the content or usage of links and references to non-Aelita Web sites or technical documentation.

No part of this document may be reproduced, stored or transmitted in any form, by any means, or for any purpose, without the express written permission of Aelita Software Corporation.

Aelita, Aelita Software, the Aelita Software Corporation logo, and all Aelita product names and slogans are either registered trademarks or trademarks of Aelita Software Corporation. Other product or company names mentioned herein may be trademarks of their respective owners.

Copyright © 1997-2003, Aelita Software Corporation. All rights reserved.

Last revised September 10, 2003

AELITA SOFTWARE CORPORATION

6500 Emerald Parkway  
Suite 400  
Dublin, Ohio 43016, USA

Phone: 614-336-9223  
1-800-263-0036  
Fax: 614-761-9620  
Email: [info@aelita.com](mailto:info@aelita.com)  
URL: [www.aelita.com](http://www.aelita.com)

---

# CONTENTS

<b>OBJECTIVE AND SCOPE</b> .....	1
<b>AUDIENCE</b> .....	2
<b>INTRODUCTION</b> .....	3
<b>MESSAGE-LEVEL EXCHANGE RECOVERY IN THE REAL WORLD</b> .....	4
<i>Deleted Items Recovery in Exchange</i> .....	4
<i>Scenarios Driving Message-level Item Recovery</i> .....	5
<i>Options for Message-level Exchange Recovery</i> .....	6
<b>OPTION 1: FULL RESTORE FROM FULL BACKUP</b> .....	7
<i>Steps, Steps and More Steps...</i> .....	7
<i>The Recovery Server</i> .....	8
<i>Pros and Cons of Full Restore from Full Backup</i> .....	9
<b>OPTION 2: FULL RESTORE FROM SNAPSHOT BACKUP</b> .....	10
<i>Impact on Recovery</i> .....	11
<i>Pros and Cons of Full Restore from Snapshot Backup</i> .....	12
<b>OPTION 3: MESSAGE-LEVEL RESTORE FROM MESSAGE-LEVEL BACKUP</b> .....	13
<i>It Seems There Is Always a Catch</i> .....	13
<i>Pros and Cons of Message-level Restore from Message-level Backup</i> .....	15
<b>OPTION 4: MESSAGE-LEVEL RESTORE FROM EXISTING BACKUPS</b> .....	16
<i>But What About My Backup</i> .....	17
<i>Addressing the Drivers With Message-level Recovery</i> .....	17
<i>Extending the Value of Message-level Recovery to the Whole Organization</i> .....	18
<i>Pros and Cons of Message-level Restore From Existing Backups</i> .....	18
<b>EXCHANGE 2003 ADDS VALUE TO THE RECOVERY PROCESS</b> .....	19
<b>SUMMARY OF BEST PRACTICES</b> .....	20
<b>ABOUT AELITA SOFTWARE CORPORATION</b> .....	21



---

## **OBJECTIVE AND SCOPE**

This document is intended to provide insight into the options for message-level Exchange recovery available in today's market. Each option is outlined in terms of its benefits and shortcomings. Based on the evaluation and comparison of these options, the document provides suggested best practices for each option, and potential scenarios for which each option may be best suited. The individual parameters of each reader's organization will dictate which option is most appropriate for them.

This document is not intended to compare the details of Exchange data backup options. Backup options will be discussed only as they relate to the ability to recover message-level items. Further, this document does not recommend one backup solution over another. The position of this document is that no matter what backup method they choose, organizations can indeed leverage their investments in these backup solutions and processes while reaping the benefits of rapid message-level Exchange restores.

---

# AUDIENCE

The intended audience for this document includes Exchange Administrators, Storage Administrators and IT Managers. Any Exchange or Storage Administrator that has gone through the pain of a day-long process to locate and restore specific Exchange content will understand and appreciate the value of the practices suggested in this document. IT management can appreciate the value of their administrators' time, and the value that their services provide to the business on a daily basis. This document offers best practices that can save both time and money for the administrator and for the organization, while extending the value of the critically important Exchange-based e-mail environment.

---

# INTRODUCTION

For decades, the telephone has been a ubiquitous form of communication. Personal and corporate interactions have depended upon telecommunications as much or more than any other medium. Today, however, one could argue that e-mail is quickly becoming a more important medium than the telephone. When people meet for the first time, e-mail addresses are exchanged as often as phone numbers. All types of business are regularly transacted via e-mail, and it is becoming the medium of choice for follow-up. Global commerce certainly has benefited from e-mail.

It's clear—e-mail has rapidly become critical to the world.

As interesting as it would be to discuss how many terabytes of data are transferred over e-mail every day and how many people have not just one but multiple e-mail accounts, this document is not about e-mail itself. Instead, it's about the significance of the commercial information and corporate intelligence that is transferred through e-mail every day.

Given the ubiquitous use of e-mail in business, an enormous amount of information is transferred through SMTP and HTTP, passed from one e-mail client to another, edited multiple times, stored for future use, and deleted—on purpose or accidentally. But why do we store e-mail? It's not so we can prove that we are able to save everything. Saving something—anything—is useless unless we can pull it out of storage and act on it in the future. The same is true for e-mail. Unless we can efficiently access the information stored in an e-mail backup, the backup itself is not very helpful.

Given the immediate and cost-conscious nature of the world today, unless we can quickly and cost-effectively search, restore and act on documents and intelligence, there really is little value in holding on to it. The return on investment is especially questionable when we realize that the cost of current processes for search and recovery can actually outweigh the benefit of the information in the e-mail being recovered.

Retrieval of complete e-mail information stores can be done quite effectively using today's available solutions. These solutions remain the best option for disaster recovery. However, the real challenge that continues to be difficult for messaging administrators is the recovery of specific content at the most granular level of a backup. Restoration of these message-level items can be cumbersome, time-consuming and costly. Efforts to remedy the situation have in some cases increased the storage requirements for the backup. But why should administrators change their backup process when what they really need to improve is their recovery process?

The focus of this paper will be that very dilemma. The paper will outline the demand for message-level recovery, discuss the options that administrators and their organizations have today, and provide recommended best practices for using each of the options. Finally, the paper will discuss recent advancements that enable the rapid recovery of message-level items without the requirement of altering the e-mail backup process. Specifically, this paper will discuss message-level recovery for Microsoft Exchange-dependent organizations.

---

# MESSAGE-LEVEL EXCHANGE RECOVERY IN THE REAL WORLD

Previously considered too time-intensive a process, the recovery of certain Exchange items was a taboo among administrators. The only occasion that warranted the required time and manual effort was a request from high atop the executive ladder. An executive may have deleted items or the executive team, legal and sometimes HR departments may have been in the middle of resolving an internal or external dispute. Fulfilling these urgent, high-level requests was a necessary evil, made worse only by the sheer time required to respond. No one enjoys making executives wait. In other, less visible scenarios, the answer was often, “Sorry Mr. End User, I can't help you.”

The reality of this scenario is that instead of accepting this as an answer moving forward, users tend to compile mountains of data in folders on their hard drives. They never delete anything. Sure, it helps with server storage, but it also hides a great deal of corporate intelligence from the rest of the organization. Given the valuable work done by all members of the organization, all users should have access to their information and be able to receive the same level of service as the executives.

## DELETED ITEMS RECOVERY IN EXCHANGE

---

Fortunately, Microsoft has helped out quite a bit in this area. The *Recover Deleted Items* feature in Exchange allows end users to restore their own message-level items with a few clicks of the mouse. Deleted items that were once gone forever are now stored on the production server for a defined period of time—seven to 30 days on average. Thank you, Microsoft. However, there are issues with this *Recover Deleted Items* feature:

- Storage space—this feature must be activated for all users on the Exchange server. Given the ability to have more and larger information stores on one Exchange server, the decision to turn on the feature could have major repercussions on storage space requirements.
- Time—this feature appropriately has a time limit. Otherwise the storage space on the production servers would be unmanageable. The problem then becomes how to recover deleted items that go beyond the time limits of the feature.

Research suggests:  
*Information from Ferris Research indicates that less than 35 percent of Exchange-based e-mail systems use the Recover Deleted Items feature for one reason or another.*

For more on features in Exchange, see

[www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/exchange/Default.asp](http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/exchange/Default.asp).

Unfortunately, solving the message-level item recovery problem is not as simple as turning on the *Recover Deleted Items* feature. Organizations' dependence on e-mail and the wealth of corporate intelligence held in e-mail drives the need to access backups for a variety of other reasons. Here are a few scenarios in order of complexity for the administrator—and coincidentally, in order of strategic importance to the business.

## SCENARIOS DRIVING MESSAGE-LEVEL ITEM RECOVERY

Expanding their reach: Originally targeted at hard copy documents, some regulations are a growing concern for businesses around the world. Regulating bodies, including the SEC, DoD and HIPAA in the US, the FSA in the UK, and the COB in France, are paying closer attention to e-mail as a medium for document storage and a source of important corporate intelligence.

- *Simple*—An end user deletes an item or multiple items by accident, or needs to retrieve information that was deleted some time ago. In the past, users were simply told no. Sometimes, this need is satisfied by enabling the *Recover Deleted Items* feature in Exchange. But, what happens when the feature is turned off or the time limit has expired?
- *Not so simple*—An end user or administrator with the appropriate permissions has accidentally deleted a public folder, and an entire team's documentation along with it. The bad news is that native tools, such as the *Exmerge Utility*, do not handle public folders well—so what do you do?
- *Complex*—HR is investigating allegations made by one employee against another and they need all e-mails to, from and about those employees. This has nothing to do with deleted items. You don't want the user to know, so you need to make sure you can restore to a location other than the originating source mailbox. How do you do this quickly and discreetly?
- *Sophisticated*—An outside audit or legal investigation involves multiple employees and their e-mail over a long period of time. You're looking for subject matter, not necessarily in known mailboxes so the task could include multiple mailboxes and multiple information stores. Legal deadlines and penalties, the risk of bad press, and potential customer confidence problems are concerns here. How do you get this massive task completed in a timely manner and still get the rest of your projects done?
- *Daunting*—A regulatory investigation involves government or industry agencies, complete with tight deadlines and stiff penalties for non-compliance. Maybe HIPAA regulations require all communication regarding a particular patient. Maybe the SEC is looking for all e-mail regarding a particular financial filing. The risks of not being able to fulfill these requests quickly include those mentioned above, but also may involve an impact on your company's market valuation if the financial community finds out. How about this for unwanted exposure?

Given the fact that these scenarios happen in business all too often, what options do organizations have at their disposal to address them—quickly and without overwhelming costs? The remainder of this paper will discuss four options for message-level Exchange recovery. The benefits and shortcomings of each, along with the recommended best practice for each option, will be covered.

## **OPTIONS FOR MESSAGE-LEVEL EXCHANGE RECOVERY**

---

- 1) Full restore from full backup—uses regular Exchange backups, but involves a cumbersome process to reach specific message-level items.
- 2) Full restore from hardware snapshot backup—uses a much faster backup process, but still requires a difficult message-level restore.
- 3) Message-level restore from message-level backup—provides a rapid restore of message-level items, but includes a high price to be paid upfront for this type of backup process.
- 4) Message-level restore from existing backups—leverages the regular Exchange backups already in place, and provides rapid restore for message-level items.

E-mail archiving solutions are evolving and can add to an organization's ability to address the above scenarios. Archiving is particularly targeted at the daunting compliance issue. These solutions rely on WORM (write once, ready many) storage models and often promise comprehensive search and granular restore capabilities. The problem is that these solutions are only helpful moving forward in time from their point of implementation. They are also often prohibitively expensive and are sometimes seen as "too much, just in case we have a problem." With that said, e-mail archiving is a topic for another paper. For more information on e-mail archiving, look to industry analysts such as Gartner [www.gartner.com](http://www.gartner.com), META Group [www.metagroup.com](http://www.metagroup.com), Ferris Research [www.ferris.com](http://www.ferris.com), and The Radicati Group [www.radicati.com](http://www.radicati.com).

---

# OPTION 1: FULL RESTORE FROM FULL BACKUP

Clearly, this is the most commonly used option for Exchange recovery, but it is also the most time-intensive process of the four options to be discussed. When the administrator is *unable* to recover items for an end user, this is probably the process being avoided. Whether native Exchange tools and processes are being used, or those of a third-party backup and recovery vendor, this traditional process for message-level recovery from a regular backup is a challenge. Why is it such a challenge? The simple answer given by most administrators is: time.

When the recovery process is successfully completed, it's very effective for complete information store disaster recovery. In fact, even with the other options available, this recovery option is the best for full disaster recovery. However, this method is less effective when it comes to message-level items, as the information store recovery must be completed *before* the message-level process can begin. This adds time and complexity.

## STEPS, STEPS AND MORE STEPS...

---

Exchange communities:  
*Microsoft TechNet at  
[www.microsoft.com/technet/newsgroups/](http://www.microsoft.com/technet/newsgroups/)  
offers some of the more  
popular Exchange  
discussion boards.  
Other site links are  
available through the  
TechNet Community  
Center / Related  
Communities pages.*

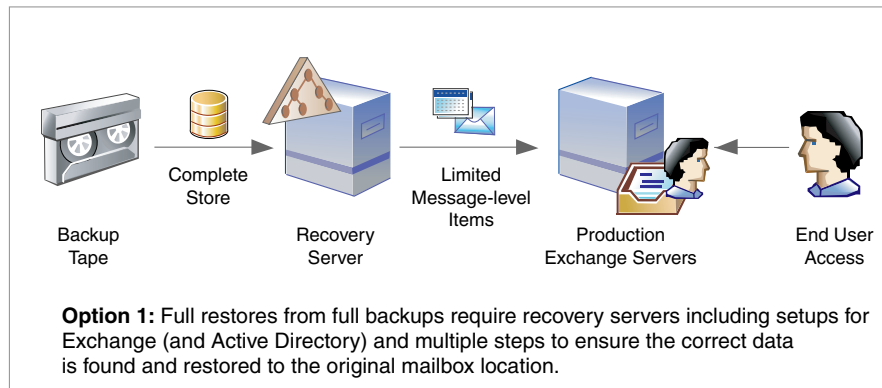
Although Microsoft and the third-party vendors have sufficiently outlined the lengthy step-by-step process, there are many bumps in the road and places in which errors can occur, making the restore process go astray. If you browse through the Exchange message boards and discussion groups, you'll find that the strings about recovery often begin with the word "help!"

There are several issues at the heart of this time-intensive process. First, as mentioned above, this process cannot access the message-level item directly without recovering the full information store first. Second, unless you are recovering from a disaster when replacement of the entire information store is possible, an alternate Exchange server (recovery server) must be used as a staging area for the message-level restore process. Third, utilities such as the native *Exmerge Utility* (known as Microsoft Mailbox Merge Wizard in Exchange 2003) are not terribly efficient when it comes to searching for specific content. The search results cannot be viewed until they are restored to a .pst file. After connecting the .pst file to a client, you can view the search results, determine which are actually needed and move the items to the appropriate mailbox.

Another challenge that impacts the overall backup and recovery situation is that most third-party products use proprietary technologies. Although vendors use standard Microsoft APIs, the connection points for the APIs in and out of the backup, the storage format, and the protocols for communicating with autoloaders are different from one vendor to another. The organization using the software is locked into using that vendor's products for their entire process. This is a fundamental reason why a recommended best practice for the backup process is to standardize on only one vendor product. However, in today's market, mergers, acquisitions and other ventures make this difficult. This, in turn, makes life difficult for the administrator performing the backup and recovery processes.

## THE RECOVERY SERVER

The alternate Exchange server, known as the recovery server, is a key element in conducting a normal message-level restore. Recovering to the live environment without using the recovery server as a staging area would overwrite any new activity that occurred in the live environment since the backup was created. Native Exchange (and the majority of third-party products) requires essentially a mirror-image of the live Exchange server to be configured in order to complete the restore.



Since this configuration information is critical to the process, be sure to make note of what is needed for the recovery server each time a backup environment changes. If the Exchange environment settings are not recorded with the backups, the appropriate recovery server settings will not be made, and the recovery process will be at a standstill. The recovery server must look like the production Exchange server from which the backup was created. The following information is required to stand up the recovery server:

- Version of Exchange including Service Packs
- Legacy ExchangeDN value
- Site name
- Organization name
- Active Directory settings (Exchange 2000 and 2003)
- Administrator privileges
- Directory backup files

If searches are being run through multiple information stores, some may have been backed up on Exchange 5.5 and some on Exchange 2000. If this is the case, the recovery server must be cleaned perfectly before changing versions. It's well known that installing 5.5 and 2000 on the same machine is asking for problems. To avoid these problems, multiple recovery servers may be needed if recovering in one of the more complex scenarios discussed earlier.

Exchange 2003 offers improvements to the native backup and recovery process via the addition of the *Recovery Storage Group* and *Volume Shadow Copy Service* features. More on this will be provided later in this paper.

## **PROS AND CONS OF FULL RESTORE FROM FULL BACKUP**

---

### **Benefits**

- Restores complete information store effectively for disaster recovery
- Uses regular Exchange backups that can be created as part of the regular full operating system backups

### **Shortcomings**

- Complete information store must be restored before filtering for message-level items
- Lengthy time requirement for this manually intensive, multi-step message-level restore process
- Recovery server setup and maintenance
- Changes in the registry

### **Best practice for usage**

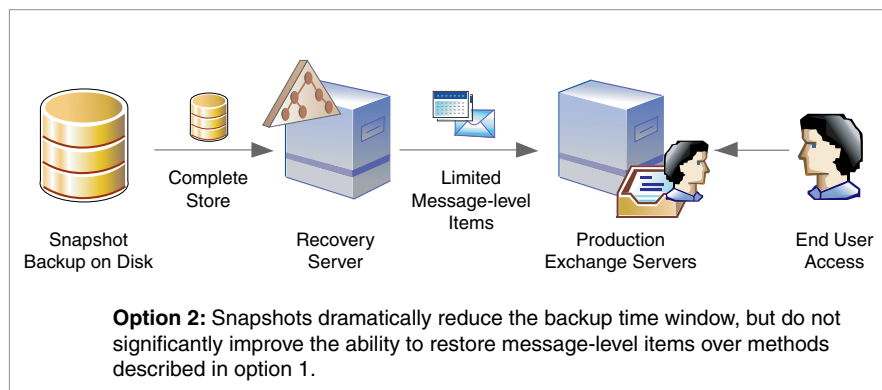
- This restore process is effective for bulk restores of complete information stores following a disaster event.

---

## OPTION 2: FULL RESTORE FROM SNAPSHOT BACKUP

Growing in popularity and affordability, snapshot technology improves the backup and recovery process on the front-end by dramatically reducing the time it takes to back up data. Rather than writing the Exchange data to tape in a linear fashion, as done in option one, snapshots can simply create a mirror (or copy) of the Exchange data and store it on disk all at once. Whether completed using intelligent storage device (ISD) hardware or snapshot-enabled software (see section on Exchange 2003 later in this document), the snapshot process can eliminate the time required to write the data to tape and provide a usable volume of the data quickly.

Snapshot technology for backups was developed to address the two conflicting requirements that e-mail administrators face every day. On one hand, users require that the e-mail service be up and running 24x7. On the other hand, the pervasive use of e-mail for nearly every aspect of daily business has created more data than can be processed in the time allocated to backup the data. If databases get too large, the backup process can impact Exchange system performance and availability. So, for the growing information store sizes of 100GB and more commonly present in enterprises today, a faster backup is needed.



Exchange snapshot backups are created in seconds, as opposed to hours. The system is able to remain available 24x7 without interruption or degraded performance. Data is captured and users are happy. But what is the impact on the recovery process?

## IMPACT ON RECOVERY

---

The process for creating a snapshot offers three recovery benefits. First, because the backup can be copied to disk, rather than written to tape, the information store is available for recovery on demand. Normal recovery begins with data being extracted from tape. The time it takes to extract the data from tape is eliminated here. However, common practice with snapshots is to store the information on disk for a period of time, and then write it to tape for longer-term storage. The benefit of having the data on disk is reduced as the information ages and the requirement for longer-term tape storage is warranted.

Second, because the snapshot process is so brief and the performance of Exchange is rarely impacted, the backup process can be completed much more frequently. The immediacy and frequency of the process reduces the number of log files created during the backup, as well as, the number accumulated between backups. Log files are essential to the recovery process because as the log files are replayed, the database is brought to a consistent state. If there are fewer log files to replay during recovery, consistency will be achieved and the restore process can commence more quickly.

Third, a typical Exchange backup is created once a day. Common practice for snapshots is to back up Exchange data as often as every hour. The availability of current backups is useful for disaster recovery of up-to-date information.

Albeit useful for recovery, a snapshot does not necessarily improve the ability to reach the granular level of Exchange data that is required in the *real world* scenarios outlined earlier.

In preparation for a normal backup, Exchange readies itself for the backup process by getting the data into a logical state. But, if the power is interrupted or the server crashes, Exchange captures data as it exists when the disaster happens. This data is in a crash-consistent state, and not necessarily a logical state. Inherent to snapshots are their crash-consistent nature. A snapshot backup does not have the steady, logical state of a regular Exchange backup. Rather than checking data pages for integrity as the pages are retrieved from disk, the snapshot only retrieves the pages—it does not check the integrity.

Additionally, implementing snapshot technologies is useful for only those backups created after the date the snapshots are first used. Discovery and leverage of existing backups still requires the cumbersome processes covered in option one above.

## **PROS AND CONS OF FULL RESTORE FROM SNAPSHOT BACKUP**

---

### **Benefits**

- Saves time in backup process
- Minimizes tape extraction delays during restore through ability to save backup on disk
- Restores complete information store quickly in the event of power loss or other disaster

### **Shortcomings**

- Recovery of message-level items is as cumbersome as with regular Exchange backups
- No benefit for recovering existing backups; only helpful moving forward
- Potentially inconsistent data
- Storing on disk instead of writing to tape can put data at risk over time
- Cost of the snapshot-enabled hardware or software

### **Best practice for usage**

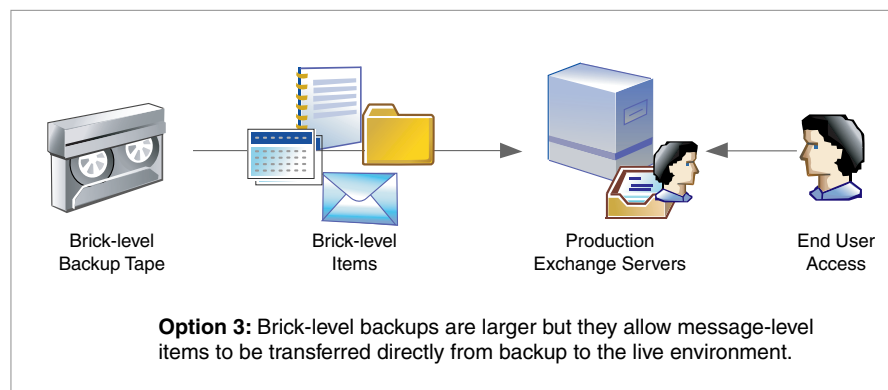
Snapshots are best used for their intended purpose—fast backups. They are great assets for capturing copies of large Exchange databases and for recovery of bulk volumes after a disaster. Bulk restore from these snapshots is effective and can save time. However, snapshots are not any more efficient at recovering message-level items than what is possible with a full restore from a regular Exchange backup as discussed in option one.

## OPTION 3: MESSAGE-LEVEL RESTORE FROM MESSAGE-LEVEL BACKUP

The brick-level backup: An easy way to understand the brick-level backup is to think of a building made of bricks. A regular backup essentially creates a picture of the building. A brick-level backup essentially creates a picture of each brick in the building. Obviously, it takes longer to take multiple pictures and those pictures take up more storage space.

Whereas snapshot technologies are focused on improving the backup process, message-level backups (commonly referred to as brick-level backups) are focused on improving the restore process, and more specifically, the granular restore process. Rather than restoring the complete information store and searching for the message-level items, the brick-level backup offers the capability to do an intelligent search for the message-level item (or items) and restore it alone. This dramatically reduces the time it takes to recover individual items.

The usual reasons for performing brick-level backups are to allow easy restoration of message-level data for regulatory, legal or emergency issues. Examples range from a corporate audit or subpoena, to the CXO deleting critical files. Brick-level backups help address the most advanced and sophisticated of the five recovery scenarios discussed earlier. Unfortunately, however, like snapshot technology, this technique is only helpful for recovering backups that occur after the solution is implemented.



### IT SEEMS THERE IS ALWAYS A CATCH

Although brick-level backups can make it very fast and convenient to restore data, they come at a higher cost than database backups for two major reasons: processing time and storage space. To understand the limitations of the brick-level backup, it's important to understand the dynamics of how it works.

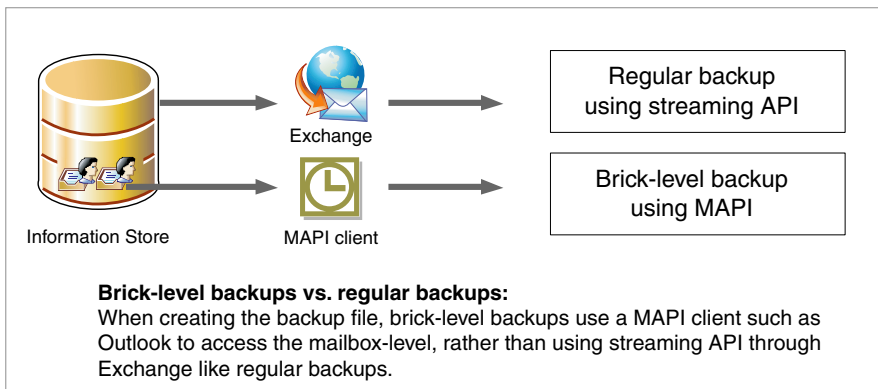
Brick-level backup solutions use MAPI calls through an Outlook or Exchange client to scan and enumerate each mailbox in the information store. This is possible because the single-instance storage model used in regular Exchange backups is broken in the brick-level backup. When this happens, each mailbox and each item within each mailbox is backed up, rather than backing up the complete database as a whole. For example, if 20 people receive the same message, a regular backup will store a single instance of the message and 20 pointers to that message. However, with the brick-level backup, a copy of the message is stored for all 20 recipients. It is easy to understand that 20 copies instead of one will consume more storage space and take more time to process.

So with this dynamic in mind, let's discuss the sacrifices organizations must make when using this powerful message-level restore tool. First, the storage space required for this backup is greater. Depending on the database, the brick-level backup can be four or more times larger. Second, since they are larger, and since MAPI calls are made (instead of high performance streaming backup API calls), the time required to create them is much longer as well. Brick-level products will vary, but it is safe to assume at least five times longer. Because the backup is larger and requires more time to create, the ability to perform the backup during slow or off business hours is in question.

The third, and least commonly known, issue is that a brick-level backup is only a supplement to the information store backup strategy. A major benefit of any backup is when it is used for disaster recovery. Exchange disaster recovery must include the complete database and its metadata. Since brick-level backups break the single-instance model that Exchange depends on, it is not suitable for a complete restore. Regular backups of the database must also be run for any users that are backed up with brick-level solutions. Storage space is impacted again because two backups of the same information are needed. Due to these challenges, most companies choose to run brick-level backups for only the VIP or executive users in the organization.

**Note:** Recent announcements from vendors such as Veritas and CA tell us that message-level backups have been improved to use the single-instance storage model for attachments. Since attachments usually take the bulk of the storage space, having only one copy of them helps to reduce the size of the backup. However, the messages associated with the attachment are still stored and recovered by breaking the single-instance model. A regular information store backup is still necessary to meet disaster recovery requirements.

Finally, MAPI-based brick-level backups can pose a challenge because of individual permission settings. Regular backups that run through Exchange and use streaming APIs access the information store as a whole and bypass the permissions set at the mailbox-level. Because brick-level backups are accessing the individual mailboxes through a MAPI client like Outlook, they are affected by individual mailbox permission settings.



If a user has restricted access to one or more of their folders, the folders will not be backed up by the brick-level process. To work around this, the Exchange administrator would need to set up a special service account that would bypass all user permission settings.

## **PROS AND CONS OF MESSAGE-LEVEL RESTORE FROM MESSAGE-LEVEL BACKUP**

---

### **Benefits**

- Message-level restore process is very fast

### **Shortcomings**

- Very long processing times for creating brick-level backups
- Regular Exchange database backups are still required for disaster recovery
- Greater storage space requirements
- No benefit for recovering existing backups; only helpful moving forward
- Special service account to bypass all permissions set of users' folders

### **Best practice for usage**

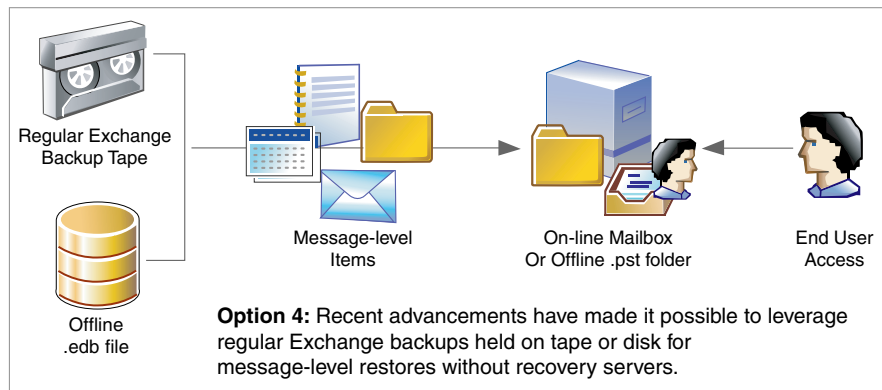
If an organization is currently using message-level backups, they are best reserved for VIP or executive users only. Exchange data from the *normal* users should be backed up using regular backup processes referenced in options one or two of this document.

## OPTION 4: MESSAGE-LEVEL RESTORE FROM EXISTING BACKUPS

The Aelita solution:  
*Linda Thacker, a systems network engineer at the United States Joint Forces Staff College, sees the value of Aelita Recovery Manager for Exchange in providing better service to the users she supports, many of whom are high-ranking military officers. "Trying to restore a mailbox is a nightmare. It sometimes takes a week to set up the recovery server and go through all the necessary steps. With Aelita I can be much more responsive. Time is critical in this job, and this product lets me recover items in 90 percent less time."*

Until recently, organizations had no choice but to use one of the options already discussed. However, with recent solution offerings from vendors like Aelita Software, the message-level recovery process can be streamlined to eliminate the time-intensive recovery servers and brick-level backups. This solution allows direct access to the backup data held in tape backups and .edb files on disk. With this access, message-level recovery can happen in much less time.

Using the regular Exchange database backup, the message-level recovery solution is able to locate and restore exactly the items required—whether the scenario involves a single message in one mailbox, or a series of items held in different mailboxes and related to a particular subject. Whether an end user needs to access the information held in a message that has already been purged from the *Recover Deleted Items* area, or the corporate auditor needs all the e-mails related to a particular client or financial filing, the process can be dramatically accelerated with this new type of tool.



Message-level recovery solutions on the market today can access MTF-compliant backup tapes. Veritas BackupExec and native Microsoft NT Backup and Exchange backups are the leading products that use this format. When working with these backup products, the message-level recovery solution can extract the backup data from tape via the MTF API. Once extracted to an .edb file on disk, the better message-level recovery solutions will bring the database to a consistent state by replaying the log files. By providing a consistent database for the search and restore process, all required data is sure to be found and usable.

Historical trivia:  
*A leader in storage management software, Veritas, actually created the MTF backup format. Written originally for Veritas backup and recovery software, the backup format now known as the Microsoft Tape Format is used natively by Microsoft and some other vendors, too.*

## **BUT WHAT ABOUT MY BACKUP**

---

Of course, there are many other backup vendors that do not use the MTF format. Instead these backups use proprietary technology developed and managed by the individual backup vendor. Many of these backup products provide the capability to restore to an alternate location, namely an offline .edb file. If this is possible, the message-level recovery solution can take over from there. If not—if the backup product forces the use of a recovery server—then the recovery process will still be rather painful.

Understandably, the existence of these proprietary technologies affords vendors many of the competitive advantages they enjoy in the backup market. But, we have also seen many of the advantages of non-proprietary or standard technologies in recent decades. As the backup vendors are able to either expose their APIs so the tapes can be accessed, or provide the capability to restore to an alternate location, their customers will be able to leverage the powerful capabilities of message-level recovery solutions.

## **ADDRESSING THE DRIVERS WITH MESSAGE-LEVEL RECOVERY**

---

Earlier in this paper, we discussed five scenarios that typically require message-level recovery. It's easy to see how message-level recovery solutions could benefit an organization facing these situations—from the simple to the daunting.

- For deleted items, it's a simple process of accessing the backup database and searching the appropriate mailbox for the deleted messages that have been purged from the *Recover Deleted Items* area.
- Public folders are like any other message-level item and can be restored quickly and easily.
- Subject matter can be searched for across particular mailboxes or an entire information store. When found in multiple mailboxes, the results can be pushed to an offline .pst file or directly into an alternate mailbox accessible by interested parties, such as HR.
- Speed is essential to satisfying legal or audit discovery requirements. Message-level recovery tools are the fastest recovery solutions available to access past backups—created with current or past versions of Exchange. Due to their ease of use, message-level recovery solutions can be given directly to the audit group so the Exchange or Storage Administrator can remain focused on other valuable projects.
- Again, speed and completeness of data is important for regulatory compliance. By bringing the database to a consistent state during the search process, a good message-level recovery tool will ensure the agencies get the data they need.

## **EXTENDING THE VALUE OF MESSAGE-LEVEL RECOVERY TO THE WHOLE ORGANIZATION**

---

Customers that use message-level recovery solutions with their existing backups clearly see the value. In addition to helping solve the original five scenarios proposed earlier in this document, the advent of streamlined message-level recovery opens new opportunities to leverage the corporate intelligence captured in existing, regular Exchange backups.

For example, consider an employee turnover scenario. Normally, when an employee leaves the business, much of the data they have hoarded on their hard drive either goes with them or is left on the hard drive, only to be wiped clean when the next user of the system is setup. Months, if not years, of corporate intelligence is lost.

However, with an easier and more administrator-friendly method of restoring message-level data from a particular mailbox, the corporate intelligence can be saved and leveraged. First of all, hard drive storage can be decreased since the reason behind massive private folders can be addressed. Secondly, since the information once stored in private folders is now stored in regular backups, the information can be restored quickly to another mailbox or an offline .pst file for another employee's use. Ramp-up for new employees can be enhanced, and the entire organization can benefit.

With any advancement, like message-level recovery solutions, it only takes a creative organization to determine what other benefits can be achieved.

## **PROS AND CONS OF MESSAGE-LEVEL RESTORE FROM EXISTING BACKUPS**

---

### **Benefits**

- Provides rapid message-level recovery
- Eliminates need to change backup process or add special backups
- Works with existing and future backups
- Restores public folders and all other message-level items including contact and calendar entries
- Enables original or alternate users to leverage information in backups

### **Shortcomings**

- Limited improvement to recovery from backups created by products that require recovery servers

### **Best practice for usage**

This is the solution for message-level recovery. Continue to use existing products, native and third-party, for disaster recovery, but this solution is a must for the granular search and restore process.

For organizations that are forced to use recovery servers, access the .edb files on the recovery server using the message-level recovery solution. The ability to search, preview, and restore items and public folders via drag-and-drop is invaluable.

---

# EXCHANGE 2003 ADDS VALUE TO THE RECOVERY PROCESS

With the release of Exchange 2003, Microsoft provides additional functionality for the backup and recovery process. Although Exchange 2003 has improvements, they are beneficial to the disaster recovery process more than to the granular message-level process. Namely, three primary functions improve disaster recovery operations:

- Recovery Storage Group
- Volume Shadow Copy service
- Microsoft Exchange Mailbox Merge Wizard (formerly known as Exmerge)

Full details of these features are available for download in various documents on the Microsoft website at [www.microsoft.com/exchange](http://www.microsoft.com/exchange).

## **Recovery Storage Group**

A new type of storage group (the Recovery Storage Group) provides a temporary location for restored Exchange data. After restoring the data to the Recovery Storage Group, you can then merge the data (using Microsoft Exchange Mailbox Merge Wizard) you need with the original mailbox store, whether that means restoring the entire information store or a few mailboxes from the information store. In many cases, this feature can replace the need for a recovery server.

Full mailboxes can be restored to the original storage group, but cannot be moved to alternate locations. Importantly, public folders cannot be restored using the Recovery Storage Group feature. The process to restore data from existing Exchange 5.5 and 2000 backups is not improved by the advancements in Exchange 2003. Be sure to consult the detailed information from Microsoft on using this feature, or you may overwrite the production content in the regular storage group.

## **Volume Shadow Copy service**

Similar to the advantages snapshots provide, the Volume Shadow Copy service creates a copy of the disk at the beginning of the full backup process. Exchange can then use this copy, rather than the live database, to produce the backup and reduce the impact on performance during the backup process. Due to API support between the new versions of Windows and Exchange, administrators should consult with Microsoft for details on using the volume shadow.

## **Microsoft Exchange Mailbox Merge Wizard**

The Mailbox Merge Wizard (Exmerge) previously available as a toolkit, is now downloadable at [www.microsoft.com/exchange/2003/updates](http://www.microsoft.com/exchange/2003/updates). With this wizard, you can move data between identical mailboxes that exist in different mailbox stores; for example, to restore a mailbox from a backup, restore the mailbox store to the Recovery Storage Group, and then use the wizard to merge the restored mailbox data with the original mailbox.

---

# SUMMARY OF BEST PRACTICES

Based on the discussion in this paper and extensive customer interaction, Aelita Software recommends the following strategies for the options covered in this document. Hopefully, these best practices serve to benefit the organization, and specifically, the IT administrator in terms of the ability to facilitate better end user service and corporate compliance with discovery and regulatory inquiries:

- The **traditional full restore** from full backup process should be used in disaster recovery scenarios where the complete information store backup can be attached to the live Exchange environment. With the advent of the *Recovery Storage Group* feature in Exchange 2003, the overhead of recovery servers is reduced, making this option less of a burden.
- **Snapshot technology** used for backups is an excellent way to capture backups frequently and quickly. An ideal application of snapshots is to marry them with the capabilities of message-level recovery so both the backup and the recovery are equally fast for the administrator and equally valuable for the end user. The log file inconsistencies of the snapshot backup can be resolved by a message-level recovery solution that replays the log files to create a consistent database.
- If being used effectively today, **brick-level backups** should continue to be used for the executive level user information stores.
- **Message-level restore from regular backup**, like that offered by Aelita Software, dramatically decreases the time and effort required to restore message-level items. By restoring message-level items directly from backup tapes or offline .edb files, the recovery server is completely eliminated and the granular search and restore process is extremely efficient. When backup vendors are able to open their APIs to allow tape access to this type of solution, the greatest benefactor will be the customer.
- For end-user-deleted items, the **Recover Deleted Items** feature in Exchange is very valuable. Storage is relatively inexpensive and provides a benefit in time and money that exceeds its cost. If turned on, the feature should be limited to a brief window of 2-5 days. Recovery of deleted items that have pushed past the time limit is best done with a message-level recovery solution discussed in option four above.
- If using a **recovery server**, be sure to document the information pertaining to the Exchange environment at the time of a backup. Document that information each time the configuration of the environment is changed. The list of pertinent information is provided earlier in this document and will make the recovery server setup process much easier. The new *Recovery Storage Group* feature in Exchange 2003 replaces the need for recovery servers in many cases.

---

# ABOUT AELITA SOFTWARE CORPORATION

Aelita Software specializes in the secure management and migration of Windows-centric IT environments. Aelita's products offer efficient control of Windows, Active Directory, Microsoft Exchange and .NET infrastructures through the secure management of migration and deployment, administration and provisioning, operations and security, and backup and recovery. These technically advanced products are engineered to reduce administrative workload, increase operational security and minimize system downtime. Aelita is a Microsoft Gold Certified Partner, as well as an HP OpenView Solution Alliance Partner. Aelita Software is headquartered in Columbus, Ohio, USA, with offices worldwide. Contact Aelita at 800.263.0036 or visit [www.aelita.com](http://www.aelita.com)

## CONTACTING AELITA SOFTWARE CORPORATION

Web:	<a href="http://www.aelita.com">www.aelita.com</a>
Technical Support:	<a href="mailto:support@aelita.com">support@aelita.com</a>
Sales:	<a href="mailto:sales@aelita.com">sales@aelita.com</a>
General Inquiries:	<a href="mailto:services@aelita.com">services@aelita.com</a>
Phone:	614-336-9223 1-800-263-0036
Fax:	614-761-9620

### **Aelita Software Corporation**

6500 Emerald Parkway  
Suite 400  
Dublin, Ohio 43016  
USA

## ABOUT THE AUTHOR

Doug Hazelman, Director of Product Management, has been with Aelita Software for three years serving in several different capacities. Currently he interfaces between prospects, customers, partners and Aelita's development team to determine the future direction and functionality of Aelita products. In addition, he has consulted on Active Directory and Exchange migration and management projects for several large companies, including Agilent Technologies, Siemens, and Unilever. Doug also collaborates on Aelita's technology publications and writes technical articles for industry journals.