


Ontrack® PowerControls™



Email Management in Today's Regulatory Environment

White Paper | 2008

KROLL ONTRACK®

Table of Contents

Why Recovering and Searching Email Archives Is Important	1
Why Recovering and Searching Email Archives Is Difficult	3
How Ontrack® PowerControls™ for Microsoft® Exchange Server Solves the Problem	5
Works with Existing Environment and Backup Process	6
A Brief Look at Ontrack PowerControls	7
The Benefits of Ontrack PowerControls	9
What to Do Next	11

Kroll Ontrack, Ontrack PowerControls and other Kroll Ontrack brand and product names referred to herein are trademarks or registered trademarks of Kroll Ontrack Inc., and/or its parent company Kroll Inc., in the United States and/or other countries. All other brand and product names are trademarks of their respective owners.

Copyright © 2008 Kroll Ontrack Inc. All Rights Reserved.

Why Recovering and Searching Email Archives Is Important

Email has become the lifeblood of business. More than any other means of communication and any business tool, companies rely on electronic messages for running virtually every aspect of their enterprises. And for many businesses, regardless of their size, the use of email means the use of mailboxes stored on Microsoft® Exchange Server.

From simple internal communications to vital sales calls to customers, to invoicing and billing and high-level decision making, email – and Microsoft Exchange Server – is involved at every level of business life. A company could not survive without email in the same way that it could not survive without telephones or electricity. And not only does email make employees and businesses more efficient, but it is also the most cost-effective means of communication as well.

Clearly, businesses worldwide are awash in email. And that email needs to be managed for business, regulatory, and legal reasons.

Business Demands for Recovering Email

Business requirements are the primary factor driving the need for email management. The main reason for this is that many companies are storing email only on the server. Kroll Ontrack, a leading provider of electronic evidence and data recovery, conducted a survey of 177 email administrators who each managed 250 or more mailboxes and found that 37 percent of companies have email stored only on the server. Because of this, storage space is at a premium and many companies are forced to set mailbox size limits. Kroll Ontrack discovered that 72 percent of companies have established mailbox quotas and for 25 percent, deleting email is the only remedy for employees when they reach a full quota. These limits force employees into deleting a great deal of email, even email that is vital to their work. Because employees delete email that they may still need, they make a greater number of requests to the IT department to restore their email – and in many cases, that email exists only as a backup. So while storing mail on the server

and setting mailbox size limits may solve the storage management needs of IT, it creates a conflict with the business needs of employees, thereby increasing the frequency of requests for restoring email messages from a backup.

Example: Kelly, an employee at XYZ Manufacturing, found a new job at another company. Before she left on her final day, she worked hard to clean up her desk and all her files. She knew that her Microsoft® Office Outlook® mailbox was always reaching its limit, so to help out she deleted all of her messages. Imagine the Exchange administrator's surprise the next week when he went in to move Kelly's messages over to her replacement's mailbox and he didn't find anything! Kelly's Inbox, Sent, Outbox, Drafts, and Deleted folders were all empty. He needed a quick and easy way to find her mailbox in one of the last backups and move all of her messages.

Regulatory Demands for Recovering Email

The need to archive and restore email is driven by regulatory demands as well. A variety of state and federal regulations require that email be kept as a normal part of doing business. For example, the Food and Drug Administration and the Security and Exchange Commission both have rules for what information must be stored and made accessible. The Health Insurance Portability and Accountability Act (HIPAA) also imposes a variety of regulations on the storage of health-related information. While the health care industry and the financial services industries are most affected by regulations, other industries are increasingly affected as well. When the Sarbanes-Oxley Act (SOX) became federal law in 2002, it impacted many companies across all industries. It established mandates and requirements for financial reporting that all publicly traded companies must adhere to, including requirements on the retention of original records such as email.

Legal Demands for Recovering Email

Email dominates communication in today's business environment. As a computer-based system, it has made this communication medium into a formal record, often containing vital pieces of information that may assist a company in tracking key events, employee behavior, and information exchanges. As electronically stored information (ESI), email messages can be valid legal documents and are governed by a variety of regulations and statutes requiring email retention. Email messages require secure storage and restoration, and often must be produced in the course of litigation. In fact, more and more companies, Qualcomm, Inc. and UBS Warburg are just two examples, have lost millions of dollars in court because they failed to adequately retain and produce email records when required to do so.

Example: Tom is a senior research scientist at XYZ Corporation and is the target of a quiet internal investigation to determine whether he is leaking company secrets to a competitor. As part of the investigation, Rachel, the Exchange administrator at his company, has been tasked with finding any messages that he has sent to suspicious email addresses (ones sent to certain domains) or any email with compromising content that he has sent to an external address. Since there is no easy way to search using this criteria, Rachel has had to painstakingly restore likely Exchange backups one at a time and do multiple searches through each one to see what she can find.

It is imperative that IT staff understand and comply with the regulations pertaining to producing email as formal records versus simply restoring email for internal business purposes. There are many software products that may assist in the recovery and production of email, but not all are compliant with the procedures required for securing email messages as documents worthy of withstanding the tests of trial.

Why Recovering and Searching Email Archives Is Difficult

Many businesses use the Exchange email server because of its full feature set and integration with Microsoft Office via the Outlook email client. But Exchange has one glaring drawback – it is very difficult to restore messages, mailboxes, and other data. To understand why, we will take a look at the database layout of Exchange.

Microsoft® Exchange Server Database Layout

A major reason why Exchange email archives are difficult to search and restore is the complexity and inflexibility of the database layout and the accompanying recovery process. Because of Exchange Server inflexibility, in many cases searching and restoring simply cannot be done.

Microsoft® Exchange Server contains both Private Information Stores (*priv.edb*) containing mailboxes, and Public Information Stores (*pub.edb*) containing public folders. The single file, *priv.edb*, is the primary database for all mailboxes. It contains mailboxes as well as email messages; however, it does not contain all email messages in all versions of Exchange Server. With Exchange Server 2000 and 2003, there is an additional file called *priv.stm* used for email as well. This companion file to *priv.edb* contains all incoming email from Exchange Internet mail service that has not yet been read. When incoming email from the Internet is read, it is added to the *priv.edb* database and deleted from *priv.stm*.

Additionally, there are several log files associated with the *priv.edb* file. These log files contain a record of all the changes that were made to the database, including all email sent and received since the last backup. When a backup is done, the log files copy their information to *priv.edb* and then the log files are emptied.

How an Exchange Backup Works

To fully understand how data can be restored, you need to first understand how Exchange is backed up. As you will see, the nature of these backups, combined with the Exchange Server database structure, makes data recovery very difficult.

There are two general types of Exchange backups: online backups and offline backups. In an online backup, the server continues to function while the backup is performed, so mail can continue to be sent and received. In an offline backup, the server is brought down during backup; therefore, mail service is disrupted. Online backups have the benefit of not disrupting the email service while offline backups have the benefit of speed. So, administrators have to balance the need for speedy backups against the downside of mail disruption when deciding whether to use online or offline backups.

Administrators also must decide whether to do a full backup or to back up an individual mailbox or individual mailboxes (a.k.a. brick-level backup). In a full backup, the entire *priv.edb*, *priv.stm* (if it exists), *pub.edb*, and associated log files are backed up. This kind of backup is ideal for disaster recovery – should a server or hard disk crash, the entire Exchange database can be restored, and so all email and mailboxes can be restored as well.

However, this kind of full backup poses a serious problem: It is an all-or-nothing approach to restoring data. You can only restore the entire database, with all mailboxes and messages. You cannot restore an individual mailbox or groups of mailboxes, and you cannot restore individual messages. Additionally, you cannot search through the backed-up mailboxes for individual messages or attachments.

There is an exception to this – an expensive, complicated method allows administrators to do a full backup and then restore individual mailboxes or messages. First, you have to build a recovery server that is a duplicate of your production Exchange server. Then, you restore the backup to the recovery server. From the recovery server you can export individual mailboxes to PST files and search through those PST files for the messages you need to recover. After all of this, you can finally copy messages back to the production server. This process is expensive and difficult to do, and the restore cannot always be accomplished. Since the recovery server must be set up exactly like the production Exchange server, if configuration information about the Exchange server has not been well-documented, you will not be able to

restore the backup to the recovery server. Additionally, buying and maintaining a recovery server requires a great deal of expense. As an alternative, you could build a recovery server only when you needed it; but, that could take an entire day, which is generally impractical in a corporate environment in which information is needed quickly.

If you have done a brick-level backup, you can restore single mailboxes, groups of mailboxes, single messages or groups of messages. However, there are drawbacks to this approach as well. Because of inefficiencies in the way data is backed up, brick-level backups take significantly more storage space than full backups and they take far more time to complete as well. For example, a server with 400 mailboxes can take about one hour to do a full online backup. The same server, doing a brick-level backup of all of those mailboxes, one at a time, can take 18 hours. Also, you cannot restore a full Exchange database using brick-level backups – for that you have to do a full backup.

Problems with Exchange Backups

This leaves administrators with difficult choices to make when deciding on a backup strategy. Should only full backups be done because they are so much less costly and resource-hungry than brick-level backups? However, with full backups there is no easy way to restore individual mailboxes and messages. This is problematic for administrators who need to find and restore individual mail messages and mailboxes. Doing brick-level backups by themselves, though, is often not practical because of the time and expense involved. And there is the additional problem that entire databases cannot be restored when doing brick-level backups.

As a result, companies are left with a less-than-ideal choice when it comes to backups. Many use a hybrid approach and do both kinds of backups – full backups and brick-level backups – on different schedules. However, because brick-level backups are so costly and time-consuming, some companies do brick-level backups on only certain mailboxes, such as those of the top executives or top managers.

How Ontrack® PowerControls™ for Microsoft® Exchange Server Solves the Problem

Ontrack® PowerControls™ for Microsoft® Exchange Server solves the problem by allowing administrators to restore individual messages, mailboxes, attachments, and even notes, contacts and tasks, from a previous full backup or from a snapshot. The software can directly read EDB files, so there is no need for doing a brick-level backup to restore individual messages and mailboxes. It lets you search across all mailboxes in an archive EDB file, rather than searching one mailbox at a time or bringing an old backup back online for analysis. You can search by a variety of criteria, including keywords, subject, date and specific users. Individual mailboxes need not be backed up because they can be restored directly from an EDB file.

Ontrack PowerControls does not require you to change your Exchange environment or your normal backup procedures. And if you happen to change your backup procedures, it will work with them as well. This is because Ontrack® PowerControls ExtractWizard™, which is included with Ontrack PowerControls, restores the database from your Exchange backup to another location. Therefore, you can restore items from legacy backups as well as from your most recent backup. Out of the box, Ontrack PowerControls ExtractWizard can extract a database from a Windows® NT Backup. Additional agents are available to allow it to extract databases from backups created with other software such as Symantec™ Backup Exec™, NetBackup™, and CA ARCserve® Backup. If you are using snapshot technology, Ontrack PowerControls for Exchange can directly read the snapshot regardless of its type (full, differential, incremental, etc.)

Works with Existing Environment and Backup Process

To better understand how Ontrack PowerControls can be used without changing your existing Exchange environment or backup process, you need an understanding of the architecture. Figure 1 shows a schematic of how it works.

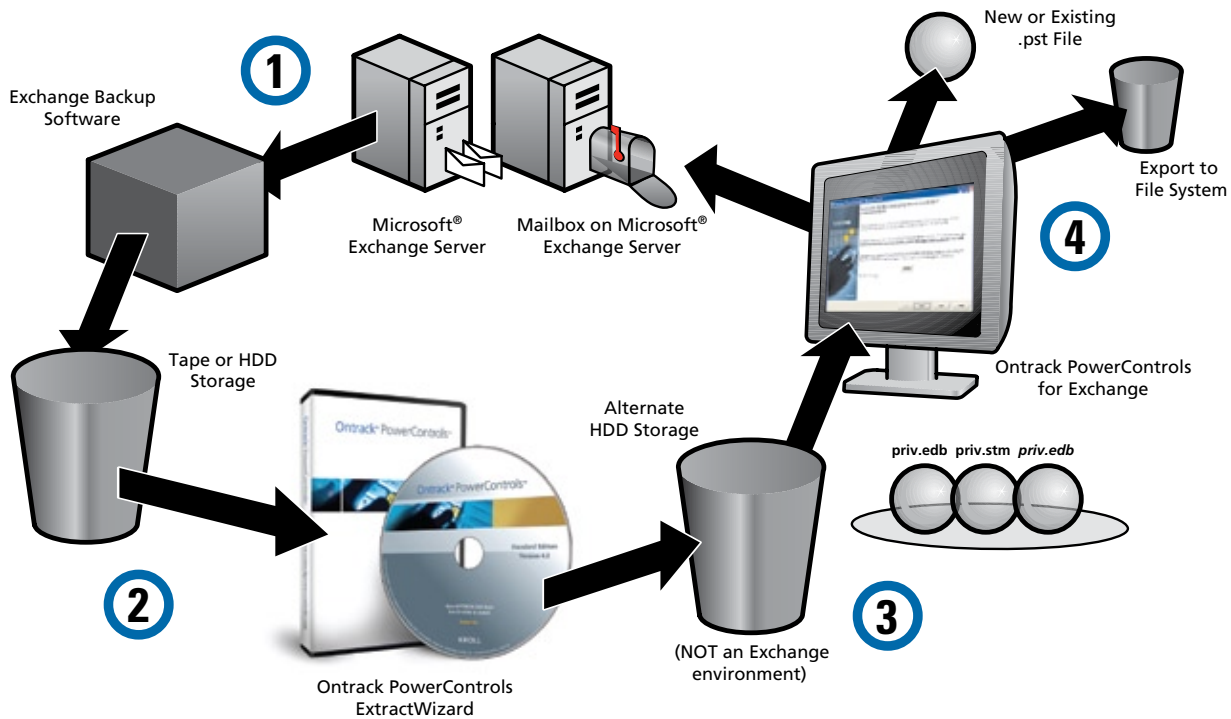


Figure 1

The key to Ontrack PowerControls for Exchange is its ability to directly read EDB files. Using it requires no change to backup procedures – rather, it works after a backup or snapshot has been completed.

- Step 1:** As normal, use backup software to back up an Exchange database and create Exchange backup sets. Or, use your snapshot software to take a point-in-time snapshot of the Exchange database.
- Step 2:** The Ontrack PowerControls ExtractWizard extracts the Information Stores from the backup and restores the database to an alternate location that is not an Exchange server. If you are using snapshots, you do not need to use Ontrack PowerControls ExtractWizard so you can skip this step.
Note: For those backup formats currently not supported by Ontrack PowerControls ExtractWizard, the Exchange backup software restores the database to an alternate location that is not an Exchange server.
- Step 3:** Ontrack PowerControls for Exchange can now be used to view and search through individual mailboxes, messages, and attachments.
- Step 4:** Ontrack PowerControls can restore single mailboxes, messages, or attachments back to the production Exchange server, or to a new or existing PST file at another location.

A Brief Look at Ontrack PowerControls for Exchange

Although it is an exceedingly powerful piece of software, Ontrack® PowerControls™ for Microsoft® Exchange Server is surprisingly simple to use. A Data Wizard runs when you start Ontrack PowerControls. Whether you need to restore an entire Information Store or just a few mailboxes or items, it walks you through selecting an Exchange database to restore from, deciding whether to restore to an Exchange server or PST file, and choosing the location of that server or file. **Figure 2** shows the wizard in action.

Note: If you do not have a target PST to which you want to restore messages or mailboxes, Ontrack PowerControls can create a new PST file for you.




Figure 2

Once you open the EDB file, it is easy to find the specific messages you want to restore. The Ontrack PowerControls search interface is much like the Microsoft® Office Outlook® Advanced Find dialog box. As you can see in **Figure 3**, you can search by keyword (including or excluding specific words). Additionally, you can add criteria to the Sent From, Sent To, and Date fields for your search. You can also selectively search in the message subject, message body, attachment file name, or attachment text.

When you complete your search, all messages matching your criteria are displayed. You can then drag and drop, or copy and paste, the located messages to the target folder. You can also restore an entire folder by using the same methods.

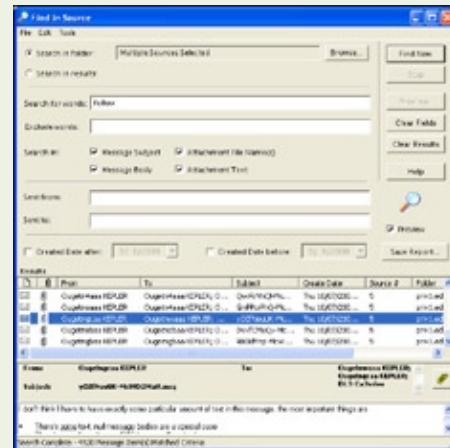


Figure 3

If you need to refine your search, Ontrack PowerControls has a “Search in results” option that allows you to conduct another search, with new criteria, that only searches the previous results.

To restore an entire mailbox, select the mailbox from the source pane and drag it to the proper place in the target. As Ontrack PowerControls for Exchange copies the mailbox, the progress and other information about the copy operation is displayed as shown in **Figure 4**. From the display, you can print a report detailing your action or you can save the report to a text file.

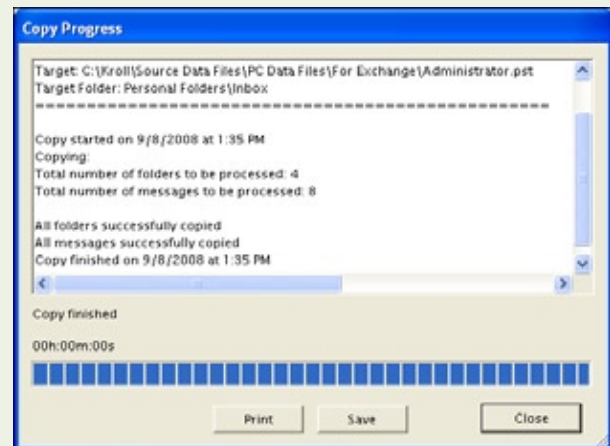


Figure 4

Ontrack PowerControls includes a powerful and useful application called Ontrack PowerControls ExtractWizard. (You use this application if you are working with full or brick-level backups; you do not need to use it if you are working with snapshots.) Most backup programs only let you restore Exchange data to the same or a duplicate server from which it was backed up. With Ontrack PowerControls ExtractWizard, you can restore Exchange data to any machine, volume, or folder that you want.

When you run Ontrack PowerControls ExtractWizard, you first tell it where the backup is located. It then searches that location for data files and displays a list of everything it finds, as you can see in **Figure 5**.

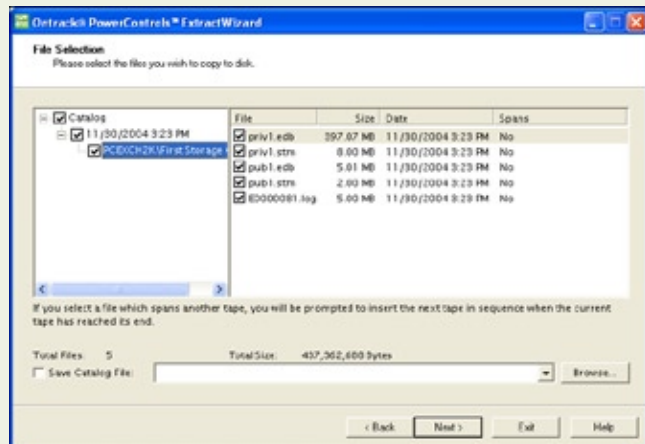


Figure 5

Choose the files you want to restore. The wizard goes to work restoring the data to your specified location. When the Wizard is done, you will see a screen like that shown in **Figure 6**.

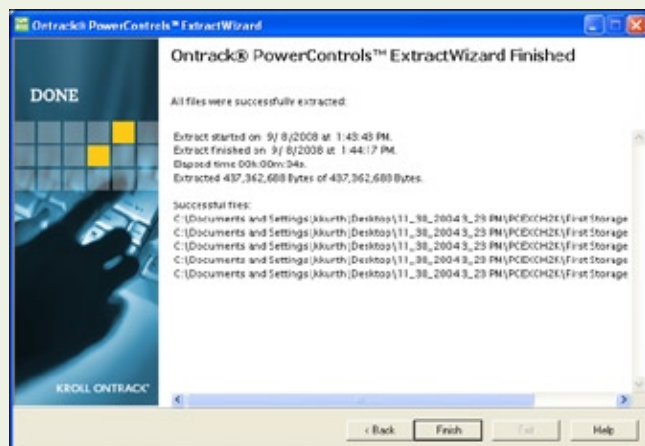


Figure 6

The Benefits of Ontrack® PowerControls™

Ontrack® PowerControls™ solves the problems that administrators face in backing up and restoring email from Exchange servers. Because it can search through and extract email from EDB files, it saves time and money, and ensures that corporations always have quick and easy access to archived email, whether needed for internal purposes, legal reasons, or for any other use.

The power of an email recovery product such as Ontrack PowerControls for Exchange expedites the search, recovery, and production of email data for use as evidence in legal cases; however, the process and use of the information must follow certain guidelines and regulations. We highly recommend that you consult with legal and technology experts such as the general counsel in a corporation, a law firm, or the legal experts at Kroll Ontrack to ensure that you are properly following all the applicable regulations.

Time Savings

Ontrack PowerControls for Exchange slashes restore time by allowing you to search for and restore single mailboxes, individual folders, or any number of messages and attachments, making brick-level backups obsolete. You can use Ontrack PowerControls to restore what you want directly to your production Exchange server or to a PST file, thereby saving enormous amounts of time and management resources. In addition, not only can it restore email messages and attachments, but also notes, contacts, calendar, and task items.

Without a brick-level backup, you would normally have to restore the entire EDB file. Now, you can use the Ontrack PowerControls advanced searching feature to look through the EDB file, find the messages or data you want to restore, and restore only those items, rather than the entire database.

Cost Savings

As we have discussed in this paper, it is possible, without Ontrack PowerControls, to recover individual mailboxes and messages using EDB files, but only by using a costly, time-consuming procedure that involves building a recovery server. However, buying

and maintaining a recovery server is an expensive proposition. In addition, the backup process itself is time-consuming and tedious, requiring substantial resources from the IT department.

So, yes, there are other options for searching for and recovering individual mailboxes and messages; however, Ontrack PowerControls for Exchange provides a much more efficient and cost effective solution than other methods of restoration for a price that is well worth the investment – typically paying for itself upon the initial use. Ontrack PowerControls for Exchange is the most solid alternative to other, more labor intensive and costly methods of recovering individual mailboxes and messages.

Major Reasons for Using Ontrack PowerControls for Microsoft® Exchange Server

There are five primary ways you can benefit from using Ontrack PowerControls for Exchange:

- Minimize the amount of storage space and the costs required to store and archive your backups. Ontrack PowerControls eliminates the need to back up mailboxes individually, thereby completely eliminating the space, cost, and time associated with performing brick-level backups.
- Greatly reduce the time required to restore an individual mailbox. Ontrack PowerControls restores mail items from a previous full backup directly into your production Exchange server or into a new or existing PST file. This eliminates the extra steps required to separately import mail back into Exchange or Microsoft® Office Outlook®.
- Minimize the time it takes to back up the Information Store. Ontrack PowerControls for Exchange completely eliminates the need to back up mailboxes individually. Many companies today perform a full Exchange backup and then run a second process to back up “Very Important Mailboxes (VIM)” individually. Ontrack PowerControls eliminates the need for this second process.
- Minimize the time it takes to locate all email matching specific criteria – keywords, specific user, subject, or date. Ontrack PowerControls has

an Advanced Find feature that can search across all mailboxes in an archive EDB file, rather than searching one mailbox at a time or bringing an old backup back online for analysis.

- Cut the time it takes to back up an individual mailbox to ZERO. Because Ontrack PowerControls for Exchange can restore mailboxes directly from the EDB file, the need for this step is eliminated!

Other Benefits

- Supports copying from a live Exchange server. Therefore, Ontrack PowerControls offers functionality to recover from archived as well as active Information Stores.
- Does not have to be installed on an Exchange server. Ontrack PowerControls is designed to run from a Windows®-based workstation and uses native Microsoft® Messaging APIs (MAPI) to communicate with the Exchange server, thus ensuring reliable and consistent operation of your server.
- Maintains integrity of the Exchange source and target data. Ontrack PowerControls does not change the metadata of the Exchange source or target files. It maintains data integrity by performing read-only operations on the source files. It will not alter data when accessing the target and maintains metadata integrity throughout the restoration process.

What to Do Next

As we have learned, granular SharePoint data search and recovery can be challenging. Ontrack PowerControls software for SharePoint offers SharePoint Administrators and IT Administrators a time-saving tool. It allows you to quickly and easily find and restore critical items without the undesirable effect of overwriting everything on your production SharePoint server or the costly expense of having a recovery SharePoint server. And, it does this without changing the item's metadata.

The best way to experience the benefits of Ontrack PowerControls for SharePoint is to try it out. Once you try it, you will see how quick and easy it is to find and restore SharePoint data. To download a free trial edition of Ontrack PowerControls, visit us at www.krollontrack.it, or for more information call us at +39 031 3525 611.

As mentioned in this paper, whether a mistake, a malicious effort, or an act of God, accidents happen and data may be lost. If you have lost business-critical data, do not panic – just use Ontrack PowerControls for SharePoint for your data recovery efforts. You can easily find any item or list in your last backup or snapshot and restore it to your production SharePoint server or to another location. If you find that the SharePoint database or SQL Server files contain severe corruption, you may want to utilize Ontrack Data Recovery, a division of Kroll Ontrack, by calling 800 44 00 33, or by visiting us at www.krollontrack.it.

Likewise, if you are faced with a request from the legal department to collect information for a discovery request, Ontrack PowerControls for SharePoint can save you time and money by allowing you to search a backup or snapshot of your SharePoint content database for the items necessary to fulfill that request. The items can then be easily copied to another SQL database for further processing and review. And, all of this is accomplished without altering the source or target files.

If you decide that you would like industry leading expert assistance on your electronic evidence data collection, processing, or production matters, contact Kroll Ontrack at 800 44 00 33, or to learn more about all Kroll Ontrack products and services, visit us at www.krollontrack.it.



KROLL ONTRACK®

Kroll Ontrack Srl
Via Lario, 1
22070 Fenegrò (CO)
Tel. + 39 031 3525 611
Fax + 39 031 3525 621
www.krollontrack.it

Copyright © 2008 Kroll Ontrack Inc.
All Rights Reserved.

Tutti gli altri marchi o nomi di prodotto
sono marchi di fabbrica o marchi di
fabbrica registrati dei rispettivi proprietari.