Condusiv™ Technologies

WHITE PAPER

Achieving Continuous Data Protection with a Recycle Bin for File Servers

Article 1:

Streamlining File Recovery Operations while Reducing IT Workload

Protecting data is a fundamental responsibility for IT operations support staff. This task is not easy. Hardware fails and data is lost or corrupted. Security breaches can result in tampering and data leaks. Employees can accidentally delete files. This guide discusses how to improve the ways you protect data, particularly for end users who access and inadvertently delete or overwrite files on network shared drives. This guide introduces the concept of a file server recycle bin to enable continuous data protection and self-service recovery, freeing IT from help desk calls to restore lost files.

This guide is organized into three articles, each covering an essential element of continuous data protection:

- Automating self-service file recovery options on file servers
- Resolving the network limits of the Windows Recycle Bin
- Implementing continuous end user data protection on file servers

Windows features on the PC, such as the Recycle Bin, are important to data protection, but they are not sufficient to provide continuous data protection or protection for network files. As a result, IT administrators lose precious time fielding lost file requests and digging through backups. Undelete®, from Condusiv Technologies, fills these gaps by offering true continuous data protection with a recycle bin for file servers. Before delving into the details, it is important to understand data recovery in the broader context of IT operations and support.

Increasing Demands on IT Operations Staff

It is hard to imagine an IT support group with too much time on their hands. The combination of changing business demands and the continuing need to monitor and maintain existing infrastructure can generate demands on IT staff faster than they can be addressed. This problem is not new.

Users have always needed support. Market conditions and business operations are constantly changing, so C-level executives are mapping out new strategies that often entail an IT component. IT support professionals and systems administrators have done their part to streamline these operations. The days when IT support staff were routinely interrupted with phone calls asking for support have given way, at least in some organizations, to help desk support systems that allow users to submit their own support requests. Rather than wait for an adverse event, such as running out of disk space or having an application attacked, systems administrators automate system health checks and run vulnerability scans to spot potential problems before they disrupt operations. Both the shift to help desk support systems and the use of automation have helped improve the efficiency of IT operations, but still more can be done. Two ways at our disposal to reduce IT workload are implementing self-service procedures and reducing dependencies on time-consuming procedures.

Streamlining with Self-Service

The use of self-service in IT support is well established. Users can reset their own passwords in many applications. The Web has become a primary resource for answering "how to" questions with common desktop applications. Users who are familiar with the Windows Recycle Bin can recover some of their own deleted files (but not all deleted files, as we will discuss in detail in the next article). You can extend the reach of self-service to include more extensive file recovery. The ability to self-recover files is particularly important for two reasons: This task is common and the recovery process is time-consuming and often requires the help of IT support staff.

Streamlining by Reducing Dependencies on Time-Consuming Tasks

By reducing dependencies on time-consuming tasks, you leave IT support staff with more time for other operations. To be sure, recovering lost files is important and can be well worth the time invested. Imagine if the final version of a sales proposal is accidentally deleted a short time before it is to be delivered to a client. The value of potentially lost business could easily exceed the time invested by IT support staff to recover that file from backup—assuming the latest version existed when the last backup was run. Not all file recovery operations will be as valuable as the lost proposal scenario, but lost files are important to the employees that lose them.

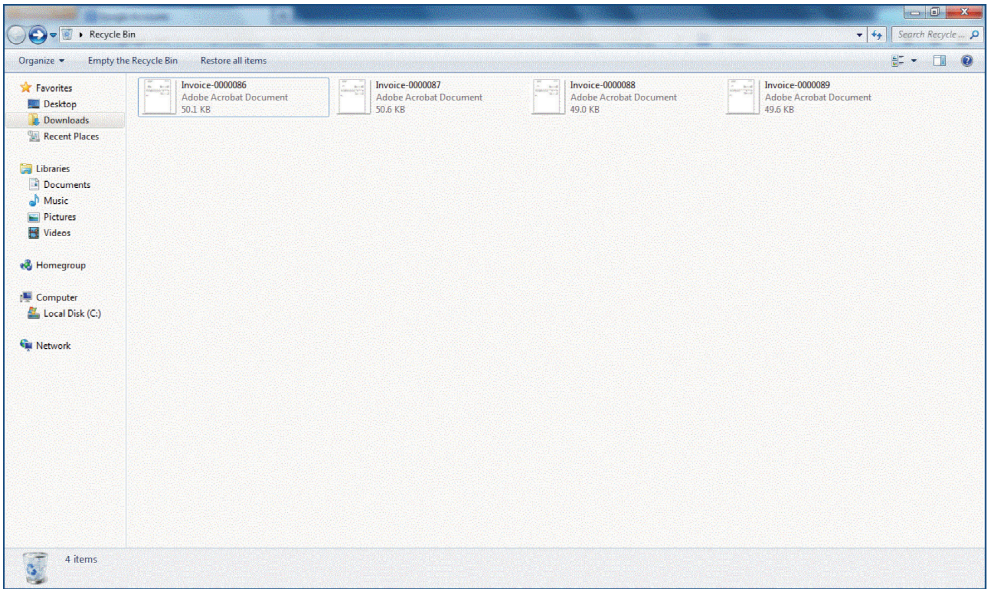

Figure 1. The Windows Recycle Bin is a first-line recovery mechanism for accidentally deleted files. Unfortunately, not all deleted files go to the Recycle Bin.

When files cannot be recovered from a user's Recycle Bin, file owners turn to IT support staff to recover the file from backups. This process can be time-consuming for all involved because the steps can include:

- The employee contacting the help desk to ask for a file to be restored
- The help desk personnel collecting enough information to precisely identify the file and version of the file to restore

- The help desk personnel needing to ensure that the person requesting the restored file actually has permission to access the file
- The help desk personnel finding time to perform the restoration operation
- The help desk determining from what backup source to recover the file(s)
- The employee having to verify the correct version of the file was actually restored

The actual sequence of events may be even more complicated. For example, if a number of files were accidentally deleted, then they may need to be restored from multiple incremental backups. The latest file may be restored before the user realizes the latest version of the file was not backed up and the user still has several hours of work ahead of him to reproduce his lost work. To avoid delays in restoring files, to mitigate the risk of lost work, and to reduce the workload on IT support staff, you can turn to self-service tools for file recovery.

File Recovery Can Be a Self-Service Operation

Many users are accustomed to checking the Recycle Bin for accidentally deleted files. What many of them may not know is that the Recycle Bin does not always capture deleted files. For example, the Recycle Bin does not capture files deleted from file servers or network drives. Those files are not available for recovery from the Recycle Bin, so IT intervention is required.

Also, previous versions of Microsoft Office files are not copied to the Recycle Bin and therefore cannot be self-recovered from the Recycle Bin. It is possible for these previous versions of a file to be captured in a backup or snapshot operation, if the file was present at the time the operation took place, in which case a systems administrator could help restore the file.

Undelete software fills the gaps in the Windows Recycle Bin functionality. Undelete allows system administrators or the users they support to recover files from a repository of saved versions of documents. Undelete operates like the Recycle Bin in the sense that it preserves a copy of a deleted file and enables recovery, but unlike the Recycle Bin, it provides this capture functionality more broadly.

As with any self-service IT support tool, you have to consider the need for security. Undelete enhances file recovery services and supports access controls comparable to those in place on the file system. For example, if a manager deleted a file named Salary_Projections.xlsx from a network share, it is captured by Undelete and stored in the deleted file repository. Only the manager, or others who have been granted access to the file at the OS level, would be able to recover that file. A curious employee from the manager's department who does not have access to the document on the network share would not have access to it in the deleted document repository.

Undelete is optimal in that it supports two modes of recovery: systems administrator-based recovery and self-service recovery. Administrator-based recovery is appropriate when dealing with sensitive information or organization policies that require a systems administrator to perform recovery operations. Self-service recovery is appropriate when the recovery system maintains Windows File Security. Self-service recovery can help streamline IT operations and reduce the workload on IT staff.

Summary

File recovery can be time-consuming for IT support staff and frustratingly slow for users. The Windows Recycle Bin is useful for recovering many, but not all, accidentally deleted files. Condusiv's Undelete is the ideal tool to provide a broader range of features and abilities to capture deleted files from network shares, as well as previous versions of files. When access controls within the file recovery application mirror the access controls of the underlying Windows OS, systems administrators can enable self-service recovery while still protecting the confidentiality and integrity of users' files.

Article 2:

Limits of the Windows Recycle Bin: Improving File Recovery Options

The Windows Recycle Bin is a well-known part of the Windows file system for the PC. For many of us, it has helped recover files that should not have been deleted in the first place. We are better off with the Recycle Bin than without it. However, as much as we appreciate the Windows Recycle Bin, we have to admit there are some significant limitations.

Windows Recycle Bin Limitations

The Windows Recycle Bin is designed to work with local files that are deleted through the Windows Explorer. This design is understandable. Windows Explorer is the primary interface to the Windows OS. Unless you are a professional programmer or a systems administrator with a particular knack for the command line, you probably spend most of your time interacting with the Windows OS through Windows Explorer or through an application such as Microsoft Word.

Limitation 1: Local Files

The fact that the Windows Recycle Bin is so closely linked to local files and Windows Explorer has become more of a problem for file recovery. We have changed the way we work with our devices and network them. Devices depend on shared network resources such as centralized storage systems. Rather than provide every desktop and laptop with hard drives large enough to meet peak demand capacity, you can configure devices with moderate-size local storage along with shared access to network storage. This method is a more cost effective and efficient way to manage storage.

Limitation 2: Deleting from the Command Line

Command line interfaces preceded graphical user interfaces (GUIs). Systems administrators and developers who work with command lines can become proficient in manipulating systems and applications with the command line. In some cases, it may be more efficient to execute a command from the command line than from the GUI. In other cases, you might need to perform a series of steps and the best solution is to use a batch script with all the commands. Working with the command line can be more efficient for many tasks, but you lose the ability to recover deleted files from the Windows Recycle Bin when you turn to the command line.

Limitation 3: Application Deleted Files

In addition to deleting files through Windows Explorer and the command line, you can work with applications that manage files. Commercial off-the-shelf software, open source tools, or custom-developed applications can support a wide range of functions that interface to the OS, including deleting files. Deleting files from an application does not entail the Windows Explorer and therefore does not support recovery of deleted files from the Recycle Bin.

Limitation 4: Previous Versions of Files

Commonly used tools such as Microsoft Word, Microsoft Excel, and Microsoft PowerPoint let users create intermediate versions of files as they are working on them. These intermediate copies are versions that are typically overwritten. This is often a space-efficient way to handle intermediate files, but it can lead to recovery problems when accidentally overwriting a file.

For example, consider working on a set of PowerPoint slides that you save every 30 minutes. Just as you are nearing the end of your work, a colleague tells you there has been a change of plans and the slides need to be substantially revised. In frustration, you delete your current set of slides and set to work on the revised version. A few minutes into your work on the new presentation, you realize that several of the slides from an earlier version of your original presentation would fit well with the new presentation. Unfortunately, that earlier version was overwritten several times. Even though you can recover the deleted file from the Windows Recycle Bin, you only have a copy of the latest version that does not contain those early slides. Had you had copies of each of your versions, you would be able to recover the slides you needed.

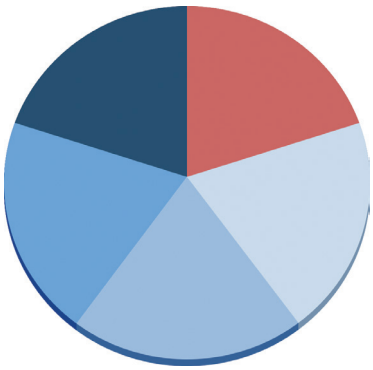

Figure 1: The Windows Recycle Bin captures files deleted from Windows Explorer, but there are other ways to delete files. Undelete is a must to protect these deleted files.

There are many ways that files can be deleted and not captured in the Windows Recycle Bin. In order to provide some level of continuous data protection, IT support staff have to find a way to compensate for the limitations of the Windows Recycle Bin.

Beyond the Recycle Bin: Options for Data Protection and File Recovery

When users cannot recover their deleted files from the Windows Recycle Bin, they may be able to recover files using one of three other options:

- Restoring from backups
- Restoring from snapshots
- Restoring quickly with Undelete

These approaches lead to varying levels of protection, so it is helpful to consider the details of each.

Restoring from Backups

Backups are commonly used for data protection and disaster recovery. Organizations have many reasons to perform backups and keep backup sets, so it is a logical assumption to think you can recover deleted files from backups. Unfortunately, there are limitations and costs associated with using backups for user file recovery.

Restoring from backups can be time-consuming for IT staff, especially when the user does not remember the name or location of the deleted file. Systems administrators may have to spend significant amounts of time searching backup sets. Similarities in files names can compound the problem because searching can return a large number of files that could potentially be the lost file.

There may be a delay before a systems administrator can restore a file. It is hard to imagine an IT support team with so little to do that they can immediately jump on a file recovery operation. Of

course, if the file or the user is a high enough priority, then the systems administrator may stop what she is doing and focus on recovering the lost file. This task, of course, just shifts the delay from the file recovery operation to someone else's support task.

Also, the latest backup may not capture recent changes to documents or files created between backups. For example, if backups are made during the night, then the latest backed up version of a file could be close to one day old. Someone working on a file all day can make so many changes that a 24-hour-old backup may be of marginal help.

The time delay and potential for not capturing the latest version of a file are classic problems in backups. These are known as recovery time objectives (RTOs) and recovery point objectives (RPOs), respectively. To reach something close to continuous data protection, you need to have near-time RPOs. To minimize the demand on IT support staff, you should have tools that have minimize the RTO as well.

Restoring from Snapshots

Snapshots of disks can preserve the state of a storage system at a given point in time. Snapshots can improve RPOs because snapshots can capture changes to files between backups if snapshots are performed during the day or at other times between backups. A disadvantage of snapshots is that they require increased storage and entail additional management overhead.

Recovering Instantly with Undelete

A third alternative for recovering files is to use Undelete as a file recovery application: consider the improvement in protection and efficiency over the Windows Recycle Bin, backups, and snapshots.

Undelete backs up versions of files as they are saved. Many users work with desktop applications and save their work intermittently. Undelete preserves all versions of such files. Sometimes users make changes to a file that should not have been made, and need to roll back to a previous version—Undelete supports that use case.

Undelete provides true continuous data protection. Back office applications, such as databases, can be configured for continuous data protection, and with Undelete, you now have exactly the same protection for users working on desktops and laptops.

In addition, Undelete captures files deleted outside of Windows Explorer. Files deleted from the command line and applications, as well as files stored on network shared drives, are completely protected.

Summary

The Windows Recycle Bin is a useful data protection tool, but it has significant limitations. Files deleted from network shared drives or from outside of Windows Explorer are not protected, nor are intermediate versions of files. Users and systems administrators have made do with the file recovery features of backups and snapshots, but Undelete is the file recovery application that provides a better alternative.

Article 3:

Continuous End User Data Protection

Businesses and other organizations often understand the need for continuous data protection for back office operations, such as databases and customer-facing applications, but may be less aware of the need for end user continuous data protection. If a bank loses data about a deposit or withdrawal, then balances will be off. A lost sales transaction can disrupt a customer's order. A lost change to a user authentication and authorization system could leave someone with excessive privileges to an application. These kinds of data losses are so potentially disruptive that large organizations take steps to provide continuous data protection. Should comparable measures be in place for end user data?

To answer the question about the need for continuous data protection for end user data, consider the following characteristics of end user data on file servers or network shares:

- Accidentally deleting files
- Accidentally overwriting files
- Maliciously deleting files
- Hardware failures and file corruption

NOTE: These recovery limitations are discussed in detail in the earlier article, "Limits of the Windows Recycle Bin: Improving File Recovery Options."

Many of us have accidentally deleted a file. Often it is not too much of a problem if we can recover the file from the Windows Recycle Bin. Unfortunately, there are a number of ways we can delete files in which they are not made recoverable by the Windows Recycle Bin.

You can also delete multiple files, for example, by deleting a folder or using a wildcard in a delete command issued on the command line.

Accidents can happen within applications as well. Such is especially the case when you use "Save As" commands and accidentally overwrite an existing file. For example, you might intend to save a document as one type, such as a .docx, but accidentally select another type, such as .pdf, and overwrite the prior version of the file.

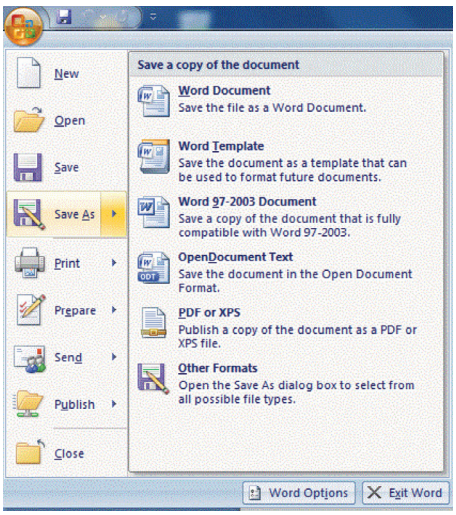

Figure 1: The Save As command can lead to accidentally overwriting existing files.

Not all file loss is accidental. Malicious attackers, from disgruntled employees to any number of targeted viruses, may delete files that appear important, randomly select files, or methodically attempt to delete all accessible files. Security measures are important to mitigate the risk of such attacks. In spite of best efforts, though, organizations can still suffer malicious attacks

that lead to data loss. In such cases, it is important to have an incident response plan in place to respond to the attack. Responses should include measures to contain the damage, assess the means by which the attack was carried out, and restore operations to their pre-attack state. Continuous data protection can help mitigate the impacts of malicious attacks.

In addition to accidentally deleting files, overwriting existing files, and malicious attacks, there is the risk of hardware and software failures. Disk drives fail. Bugs in applications can corrupt data. Laptops can be dropped and broken. It is difficult to anticipate the specifics of these kinds of failures, but it is prudent to plan for them.

Limits of Backups and Snapshots

Backups and snapshots are useful for data protection, but they are typically not a comprehensive solution for continuous data protection of end user files. For example, backups do not recover changes made since the last backup. Work performed between the time of the last backup and the time of file loss is not protected by the backup.

Figure 2: The work performed between the time of the last backup and the time of data loss is not protected by backups.

Snapshots can improve data protection between backups. With snapshots, copies of data on persistent storage devices are made at intervals between backups. This setup improves the RPO but does not meet the ultimate goal of continuous data protection. Furthermore, snapshots entail significant storage overhead. Managing snapshots and ensuring the right balance of space utilization and data protection can add to systems administrators' workloads.

Undelete: True End-User Continuous Data Protection

Undelete is a file recovery application that provides true continuous data protection with many key features:

- Undelete allows for self-service recovery of files deleted from a file server or network share. End users do not have to solicit the help of systems administrators for simple recovery operations. This setup reduces the workload on administrators and can help speed the time to recovery for end users.
- Undelete offers the ability to capture versions of files as they are saved rather than just the last version of a file.
- Undelete compensates for the limitations of the Windows Recycle Bin by protecting files deleted at the command line or in applications, while also protecting files deleted from network file shares.

- Undelete preserves file system security controls in the file recovery applications. Users do not have access to files in the recovery repository if they do not have access to that file in the file system.

This combination of features is a benefit to both end user and IT support staff. The former have better control over the time it takes to recover deleted files while the latter are left with more time for higher-priority issues.

Summary

It is possible to offer end users continuous data protection. At the same time, those users can gain greater control over the recovery process with the use of self-service recovery features for all local and network files. IT administrators can still maintain control over file security while reducing their overall workload when it comes to restoring files. Continuous data protection provides for better RPOs and reduces the risk of losing substantial amounts of work. The Windows Recycle Bin has worked well for some cases of deleted files but not all, especially in large environments that normally use centralized storage. Condusiv's Undelete is the ideal tool to fill those gaps.

Condusiv Technologies
Condusiv is a virtual company
Visit Contact Us for addresses and additional phone numbers.
www.condusiv.com