Stand Guard Against Accidental File Erasure

By Joel Shore

Sooner or later, disasters happen to every computer user. They delete important files, on purpose or accidentally. Or they update a document and save it, overwriting the original version. The common practice of using an existing document, spreadsheet, or presentation as a starting point for a new one often ends in catastrophe when the user forgets to save the changes under a new file name.

Careers end on such errors. Or worse. With the strict document-retention requirements imposed by the Sarbanes-Oxley Act of 2002, corporations that destroy records risk severe fines; employees could even wind up in prison.

Defining the problem is simple: files that should not be deleted often are, on individual PCs and file servers. The immediate challenge is recovering them. Dealing with the larger, long-term issue of prevention and retention on a corporate-wide scale is more complex.

Products that attempt to recover deleted files have been around as long as personal computers themselves, but have a history of delivering mixed results. Today, Microsoft® Windows® itself offers the recycle bin for just this purpose. The shortcoming is that these solutions, by definition, are after-the-fact fixes. None represents a preventative solution. They cannot resurrect a file that has been overwritten, and they make no attempt to archive the many revisions that a typical document goes through in its lifetime. Stronger measures are needed. Fortunately, such a solution is available today.

The Unseen Expense

The mathematics of accidental file erasure is alarming. A PC user is likely to spend an average of one hour in a frantic effort to recover the file (or files, or an entire directory) before turning to the help desk. Just two occurrences per day – a conservative estimate for a corporate environment – translates to a minimum annual productivity loss of 520 hours, nearly 14 40-hour weeks. Office colleagues, in their attempts to help, add to lost productivity and are more likely to hurt, not help, any chance of success.

In a corporate environment, just two accidental file deletions per day can cost 520 hours in annual productivity loss and $7,800 in IT staff time

Once the IT department gets involved, costs add up quickly. With an IT technician earning $30 an hour, a single 30-minute venture to locate and restore a file from a backup tape – if it's there at all – costs $15. While $15 does not seem like much, in a corporate environment, repeating that process twice a day costs $7,800 over a full year.) Tied up for 260 hours, the total quantifiable cost for user and IT is nearly 21 work weeks, the equivalent of five months from one full-time employee.

Clearly, recovering a deleted file from a local hard drive or perhaps from a prior day's backup tapes is an expensive, time-consuming, productivity-robbing process. That's if it can be done at all. And should a crucial file be unrecoverable, the cost to the business itself could be incalculable.

But there's more.

The Long Arm of the Law

In the course of a workday, workers create, modify and delete Microsoft Word, Excel and PowerPoint documents. As changes are made, little regard is paid to archiving original versions of these documents. Files no longer needed are either deleted or overwritten.

Under the Sarbanes-Oxley Act of 2002, deleting a file could lead to serious legal consequences

That's no longer good enough. In response to the multibillion-dollar misdeeds of several major

corporations, Congress responded, enacting the Sarbanes-Oxley Act of 2002. It changed the rules for document retention within public corporations and imposes severe penalties for offenders. Sections 802 and 1102 of the Act amend the federal obstruction of justice statute, greatly increasing the penalties for the criminal destruction, alteration, and falsification of records in certain circumstances.

Under Section 802, anyone who knowingly alters, destroys, mutilates, conceals, covers up, falsifies, or makes a false entry in a document with intent to impede, obstruct or influence the investigation or administration of any matter within the jurisdiction of a federal department or agency or any bankruptcy case is subject to a fine and imprisonment for up to 20 years. Similarly, under Section 1102, anyone who corruptly alters, destroys, mutilates, or conceals a record or document with intent to impair its integrity or availability for use in an official proceeding is subject to the same penalty.

Should you wait for the government to come swooping down on your company before implementing a document-retention plan? Certainly not. The destruction of files in anticipation of a future inquiry is not a good idea; it pays to plan ahead. Of course, cooking the books is not exactly good business either.

To avoid potential legal exposure, every publicly traded corporation should develop a document-retention policy and disseminate it to all appropriate employees. An increasingly popular way to do this is by posting the policy on the internal corporate intranet. Smaller organizations need not worry about potential legal exposure under Sarbanes-Oxley, but implementation of modest document-retention plans is good business.

Software that automates and ensures the safe retention of documents – current and prior versions – is essential.

Recycle Bin Woes

Windows offers the Recycle Bin as a safety net, but its limited functionality leaves users dangerously unprotected

A corporate retention strategy looks at the big picture. What such a plan is likely to miss is the workaday fine details down in the trenches, the often unintentional actions taken by individual rank-and-file workers who spend hours at a computer every day. Put another, less-flattering way, corporate executives may erase files because they have something to hide; an average employee deletes or overwrites an individual file by accident, then desperately tries to get it back.

To deal with file erasure at the individual PC level, Windows offers a safety net, the recycle bin. The idea is that deleted files are held here, for a while, in a sort of electronic limbo, allowing a user to recover a file before it disappears for good. Though an excellent idea, this net is like any other, mostly large holes stitched together with fine thread. It contains serious shortcomings that limit its usefulness.

By default, Windows allocates roughly 10 percent of a PC's available local hard drive space for use by the recycle bin. Once the bin becomes filled, Windows simply purges the oldest files to make room for those newly deleted. Depending on the available space, deleted files might linger for weeks or only days. This limitation manifests itself in the tens of millions of otherwise healthy, still-in-service corporate and home PCs with hard-drive capacities of 20 GB or less.

The recycle bin is often quite small. In a system with a 20 GB hard drive, the recycle bin is less than 2 GB in size. (Less, because the bin is based on available drive space, not total drive capacity.) This space issue can lead to disaster for especially large files. Since any deleted file larger than that cannot fit in the bin, Windows instead simply purges it. Though the file is history, as far as the recycle bin is concerned, it may still be recoverable if the right tool is used, and used before that location on the hard drive is written over with other data. Conventional wisdom would suggest that locally stored files of this size simply do not exist, but such thinking is naïve. The data file from a salesperson's contact-management application can easily grow to this size over a period of years. And database tables, almost always the result of ad hoc departmental applications created without the knowledge of the IT staff, can grow equally large.

Network Exposure

The Windows recycle bin offers no protection for files deleted from the server by network clients

To gain greater control over user's files, it is increasingly common for IT departments to configure users' PCs to save files to a network directory, rather than the computer's local hard drive. Good reasons for doing so are plentiful, yet this scheme can further harm the already slim chances of recovering a file. The reason is simple: items deleted from a network drive are permanently deleted; they are not sent to a Windows recycle bin.

The Volume Shadow Copy Service (known as VSS), a component included in Windows Server 2003, attempts to deal with network file issues. Like the recycle bin itself, VSS performs a specific function and was not designed to resurrect deleted files.

VSS is a snapshot tool. As such, it creates a point-in-time copy – a snapshot – of a network volume. That can be a problem. Changes to existing files or newly created files are not backed up until the next snapshot occurs. Files that are overwritten by a newer version are gone for good. In addition, not all applications are VSS compliant. Applications that do not contain code supporting VSS still have their files backed up, but the relational integrity of databases spanning multiple volumes may not be preserved.

And it gets complicated. A complete VSS environment consists of a backup application (requestor), the business application creating the data (writer), and storage snapshot technology (provider). These elements must be integrated with the VSS framework and must reside on the same computer.

VSS certainly works as advertised, but, like any backup system, it can preserve only that which exists. Files that have been erased or overwritten simply aren't available for backing up. And for that reason, a separate, powerful tool that focuses exclusively on deletion protection and version control is essential.

A Better Way

Diskeeper Corporation's Undelete 5 is a proactive solution that offers complete protection from accidentally deleted files, as well as older versions of Microsoft Office files

There is no shortage of products that tackle various aspects of file recovery. Most of these are after-the-fact solutions that help a desperate worker in an attempt to recover an already-deleted file. It's akin to, as the adage goes, shutting the barn door after the horse has escaped. A more advanced approach to the problem is to not allow such a situation to occur in the first place. After all, preventative medicine is far better than an ambulance ride to the emergency room.

A complete solution for deletion protection must encompass several key aspects. Foremost, it needs to do a more complete job than the Windows recycle bin, capturing any deleted file, regardless of size or location. It must provide automatic version protection, saving multiple copies of Word, Excel and PowerPoint documents as they change throughout their lifetime, and doing so without any user action. It must be network savvy, providing protection for users that save their work to a network directory. From the IT perspective, a proper solution should have the ability to be push installed over the network to individual systems and allow the entire environment to be managed centrally.

Diskeeper Corporation developed Undelete® 5 to deal with these circumstances. Undelete replaces the Windows recycle bin with its own more powerful "Recovery Bin" that intercepts all deleted files, regardless of their size, their location, how they were deleted, or who did so. The recovery bin employs a Windows Explorer-like interface for navigation and a search capability. By right-clicking on the file and selecting Recover, the file is brought back.

It's a fact of life that the needs of individual users and network administrators are different. Recognizing this, Diskeeper Corporation opted to tailor separate editions of Undelete 5 to their divergent needs. In fact, four versions were created. The Server Edition protects server files, including those deleted by network users from their own PCs. It has two cousins. One is the Professional Edition; it allows users to protect locally-stored files and also to recover files from Undelete Server's recovery bins. The other is the Desktop Client, which allows connected workstations to recover files from Undelete 5 Server recovery bins. It's especially suited for

environments where users cannot store anything to their local hard drive. Finally, there is a Home Edition, which allows recovery of locally stored files.

The Server, Professional, and Client editions can access recovery bins on remote systems running Undelete Server Edition. This fact alone justifies the cost of acquisition: The torturous task of searching through a server's backup tapes to recover files deleted by a network user is no longer necessary. Primary and extended partitions, volume sets, mirror sets, and RAID arrays all are supported. All editions are compatible with the latest service packs for Windows XP and Windows Server 2003™.

Protection is provided that guards against losses that occur when a changed Word, Excel and PowerPoint document is saved under the same name, overwriting the prior version. Undelete 5 implements version protection, allowing users to restore these overwritten files. In use, users need only click on the file and select the View Versions option. At that point, version restoration is available.

To simplify deployment and maximize valuable IT resources, Undelete 5 incorporates a push install feature. With it, the Professional, Server, and Desktop Client editions can be installed to selected workstations and servers throughout the network. The elimination of hands-on visits to each system is a major factor in cost justification.

Diskeeper Corporation conceived of Undelete as a "Set It and Forget It"® file recovery system. (In fact, it is Diskeeper Corporation, not the maker of a rotisserie oven advertised on late-night TV, that holds a trademark to that very phrase) To do so, Undelete operates as a Windows service that runs in the background. To assure complete protection, the service, which consumes negligible system resources, runs continuously, completely transparent to the user.

Conclusion

Preventing users from deleting or overwriting files is a nice idea, but one that isn't possible. It happens frequently, causing anguish for the user, aggravation for the IT department, and possible financial and legal harm for the business. To deal with these vagaries of human nature, the use of software that backs up files or attempts to recover them is logical. Unfortunately, these products are often reactive in nature. They cannot recover files across network directories nor do they provide protection for existing files overwritten by a newer version. Furthermore, backups are intended for disaster recovery from a catastrophic systems failure, not as a method for file-by-file undeletion.

By giving administrators a comprehensive file recovery and version-protection solution with instant recovery, Undelete 5 provides cost savings—and peace of mind

Undelete 5 was built with these challenges in mind. Its ability to recover files of any size, from any local or network location, and to provide comprehensive version protection and restoration services, sets it apart from other products currently available. These capabilities are certain to put users' minds at ease while allowing precious IT resources to be redirected to core-business projects.

This report was sponsored by Diskeeper Corporation, however, the sponsor had no input into the content of this report, conclusions reached, or opinions expressed by Reference Guide.