How To Prevent Data Loss

Data loss can be devastating to any business, but recovery is much easier if a contingency plan is formulated before the event, rather than after.

By Dave Cook
Technical Journalist

Forget about motherboards, processors, memory modules and graphics cards - the most important part of any computer system is the data it holds. Should data loss occur, it could take ages to re-install those bloated applications, let alone any vital files and documents. But time is money, and in today's business environment the information and data a company stores is often the basis of its competitive advantage. With the temporary loss of such data, a company can quickly lose its cutting edge. To lose it permanently will almost certainly bring about severe disruption; so much disruption that many companies will never recover from such a loss.

While disk drives are more reliable than they have ever been, disk failure is not the only reason for misfortune to strike. Apart from fire, flood, and other "acts of God", the system administrator has to look out for a multitude of potentially disastrous possibilities, including the loss of data from computer virus, human error, theft, and deliberate vandalism.

Have A Recovery Plan

How quickly a company gets up and running after such a "disaster" depends largely on the precautions it has taken beforehand. After all, it is far better to formulate a recovery plan before the event, rather than later. The type of recovery plan chosen will depend not only on the level of failure or downtime the company is prepared to accept, but on how much money it is willing to spend on a recovery strategy. A real-time fault-tolerant system - preferably achieved by remote server clustering - should ensure continued operation after failure occurs. If a server were to go down, a real-time fault-tolerant system would automatically switch to another server or system. No data would be lost, and the least amount of disruption would occur.

But real-time fault tolerance is not cheap, and neither is it always perfect. Typically, when a file is deleted from a real-time fault-tolerant server, the server cluster also deletes the file. Recovery, therefore, is hardly "real-time" because it usually takes several minutes to restore the file from its backup set. Moreover, the file may not be identical to the lost file; it could be hours, or perhaps even days old. A short-time fault-tolerance strategy is a far cheaper alternative. This type of strategy is best for businesses that can survive a downtime of around two or three hours without grinding to a halt. During downtime all company files and records are handled manually or moved to another server or workstation until the problem is fixed.

Protection

At least Windows NT offers its own degree of fault tolerance. Some of these protective features work transparently, while others, like the Emergency Repair Disk and Last Known Good configuration utilities, are designed to get the system up and running again with the minimum of downtime. The Emergency Repair Disk is basically an NT-formatted floppy containing the files found in the %SYSTEMROOT%\REPAIR folder. Although the first disk will have been created during the initial NT installation, it can be updated thereafter using the RDISK utility from the command prompt. It is important to note, however, that for reasons of size - this could be several megabytes on certain systems - the RDISK utility does not update the SAM and SECURITY hives. Employing RDISK with

the /S switch will copy these hives to the repair folder, but they should be backed up as part of the normal backup procedure as well.

Typically, the Emergency Repair Disk contains enough registry and file setup information to return the system to a bootable state. However, the disk must be kept up to date. It should be updated before any major changes are made to the system, and then updated again after the changes have been made and the system is fully operational. The Last Known Good feature can also be a quick recovery tool. It works because most successful NT boots are cloned to the LastKnown-Good entry in the system registry. Thus, hitting the space bar when NT boots will invoke the Last Known Good configuration, effectively ignoring the new driver or system settings responsible for having caused the problem.

More recovery options can be found in the System Properties Startup / Shutdown tab, accessed by clicking on the Control Panel system applet. For example, the system can be set to reboot immediately after a crash and without the need for any manual intervention. This is useful should the crash occur at a time when the server is not being monitored.

RAID

Hard drive integrity is, of course, essential to any system. Hardware RAID in particular can offer superb protection from disk failure. This is because data is spread or striped across a set of disks, improving throughput and protecting the data held on the array against the failure of any individual disk. There are six different ways to configure an array. Levels 0, 1, 3 and 5 are the most popular, though each level contains benefits and drawbacks. However, even the best RAID is not completely infallible. Two disks could fail together, for instance, or the controller could develop a fault. Nevertheless, a properly configured RAID should be considered an invaluable asset, especially for the larger network.

UPS

Data protection has many facets, of course, and when looking at the larger picture it is easy to forget the one thing that most of us take for granted - the power supply. Indeed, many experts believe that external power failures and surges account for up to 40% of all system failures resulting in data loss. The best way of avoiding the problem is to install an uninterruptible power supply - commonly known as a UPS. Basically, the average UPS is a box containing lead-acid batteries, complete with an inverter to convert direct current from the batteries into an alternating current.

The most basic UPS will protect a system from any sudden and unpredictable power surges. More expensive units are designed to cope with a total blackout. In such cases, a good UPS will keep a system up and running long enough to administer orderly shutdown procedures. For networks that need to be up every minute of every single day, even more powerful UPS units are available. Typically, these UPS units will run for several hours, long enough to give an emergency generator time to kick in.

Backup Strategy

As always, a solid backup strategy is vital to any recovery plan and, for sheer convenience, tape is still hard to beat. However, because today's hard disk capacities are huge it would be extremely time-consuming to perform a full backup more often than is necessary. Some small businesses, for instance, may require a full backup only once a week, with differential or incremental backups making do in between.

A differential backup copies files that have changed since the last full backup. Although this type of backup is not as time-consuming as creating a full backup, both the full and differential backup tapes are required to fully restore a system. Incremental backups are quicker still, since they contain only the files that have changed since the last backup of any description. When it comes to the restore process, however, incremental backups are not quite so convenient because several backup tapes may be needed to get the system back to normal.

The most popular backup strategy is based on a three-generation system, with 21 tapes used to back up data for 12 months. This strategy is commonly known as

Be Prepared

Getting a system up and running after a failure can be much easier if a recovery kit is made readily available. An example of a Windows NT recovery kit is as follows:

- A DOS boot disk.
- An NT boot disk.
- A recently updated Emergency Repair Disk.
- A recent full backup.
- If a full backup is not available, a backup of the registry using the REGBACK utility.
- All relevant documentation concerning the installation of the server and its applications.

the grandfather, father and son set. Typically, a full backup is taken at the beginning of each month (grandfather), another full backup at the start of each week (father), with differential or incremental backups in between (son). By following this approach it is possible to create day-to-day backups of the current week, week-to-week backups of the current month, and month-to-month backups of the current year. Backups should always be verified and the restore process tested frequently. At the end of the year, the month and year tapes can be archived before starting afresh with new ones.

At the very least, backups should be stored in a fireproof safe. Even so, this is not ideal. After all, what good is the latest backup if it cannot be accessed? Typically, a fire-ravaged building could be so badly damaged that it is off-limits to everyone, and it could be days or even weeks before access is granted. Hence the best solution is to keep backup sets safely stored offsite. One way to achieve this is to back up crucial files across the Internet, and a growing number of providers now offer this useful alternative. For example, @Backup (**www.backup.com**) provides a 30-day free trial and charges a flat US$99 per year for 100 MB of storage.

Disk Imaging

As an extra form of insurance, utilities that transfer the entire contents of a disk - or specified folders - over to another drive are also worth considering. PowerQuest's Drive Image Pro 3.0 (**www.powerquest.com**), for instance, creates an exact image of a hard drive or a hard drive partition. This image can then be used to deploy multiple Windows workstations, upgrade existing workstations, or back up and restore hard drives. The utility is primarily designed to protect data from a major hard disk crash, but since it supports selected file restore, it can also be used as a belt-and-braces approach to backups. At its highest compression ratio, Drive Image Pro offers around 50% compression; this means that, in theory, a 22 GB drive that's three-quarters full could have its entire image saved to an empty 10 GB drive.

Virus Threats

Thanks to the growth of the Internet and our increased reliance on email, the threat of losing data from viral infection is now greater than it has ever been. There are, of course, many different types of virus, but it is the relatively recent threat of the macro virus, sent as an attachment to email, which is presently causing

"Many experts believe that external power failures and surges account for up to 40% of all system failures resulting in data loss. The best way of avoiding the problem is to install an uninterruptible power supply."

Prevention Is Best

Today's hard disks should provide many years of trouble-free service. They rarely fail, but when they do the results can be catastrophic. Therefore it pays to be aware of potential problems before they become a real threat.

- Back up data regularly. Store backup sets in a fireproof safe, and off-site wherever possible.
- Use anti-virus software and maintain regular updates.
- To help safeguard against boot sector infections, disable the "boot from floppy" option in the system BIOS.
- Store important files in a single location such as the My Documents folder. Then use subfolders to help organise data by project or category. Old data can be moved across to a specially created Archives folder.
- NT users should maintain an up-to-date Emergency Repair Disk.
- Format new NT volumes as NTFS. The NTFS format includes transaction-logging capabilities to help prevent data errors.
- Train users to report any unusual noises immediately. Tapping, clicking or humming sounds are early signs of disk or controller failure.
- Avoid excessive heat. The latest hard drives spin almost twice as fast as their older counterparts, so make sure the system fan is up to the task. Consider fitting an additional fan.
- Store an emergency installation of the OS on a different drive.
- For optimal server setup, use mirrored boot drives with all data maintained on a hardware RAID 5 system. Depending on the nature of the data involved, consider RAID 5+0 for increased speed.

most concern. At least most of the latest anti-virus products can examine attachments as they sit in the post office. Other applications can work in conjunction with Internet firewalls to search incoming mail for viruses. Alternatively, configuring the system to save all attachments to a "safe haven" can easily solve the problem. This could be to another drive, for instance, which is set to automatically scan for viruses every time a file is written to it.

Using anti-virus software is every bit as important as keeping regular backups, and the anti-virus software must be updated regularly. Meanwhile, the whole workforce should receive regular anti-virus training and be reminded constantly of the dangers of virus infection.

Data Recovery Tools

Whatever the cause, it is just as well that data loss is not always on a large scale. When it comes to the occasional file deleted by accident, there are few utilities as useful as Executive Software's Undelete for Windows NT. Available as a free download (**www.execsoft.com**), the utility is part of a fully featured Undelete package that changes the Recycle Bin to a Recover Bin, making it easier to retrieve files no matter how they have been deleted.

A number of manufacturers provide do-it-yourself solutions in the form of data recovery software. But utilities such as these should be treated with caution since, in the wrong hands, they can be responsible for losing data permanently. Not so Search & Rescue, the enterprise version of PowerQuest's Lost & Found data recovery application (**www.powerquest.com**). PowerQuest claims that as long as the disk is still spinning, Search & Rescue can locate and recover almost any file, anywhere on the disk. Even accidentally formatted disks may be cajoled into surrendering lost data, though only as long as the data hasn't been overwritten first. Importantly, no additional hard disk damage or data loss can occur because Search & Rescue never writes to the impaired disk. Once Search & Rescue locates the data, it is left to the user to decide where to transfer and restore the recovered files to.

Recovery Services

Recovery utilities often offer a quick and cheap solution to data loss problems, but they are not always the best option. When a company's very survival is at stake, the common sense approach is to call in a data recovery service sooner rather than later. Professional recovery experts can recover data from virtually every situation imaginable, and without voiding equipment warranties. Indeed, many of these organisations claim that around 95% of all inaccessible data can be recovered. Be aware that most companies offering this type of service advise users against the do-it-yourself approach, claiming that the use of data recovery utilities may decrease the chances of a successful retrieval. Others, such as Ontrack Data International (**www.ontrack.com**), a company with a long and established track record, actually provide both DIY and remote solutions as well as in-lab capabilities.

Conclusion

Unfortunately there is no such thing as a single, foolproof way to protect systems from data loss. But with a strict backup regime in place and a carefully thought out disaster contingency plan, the threat of losing data permanently can be reduced to almost zero.

PCNA

Further Information

Data recovery information and services
www.datarecoverygroup.-com

Data recovery FAQs
www.cbltech.com

Daily backups over the Internet
www.netstore.net

Crash protection software and anti-virus products
www.symantec.com

RAID technology information
www.rising-edge.com

RAID 1 IDE disk mirroring
www.arcoide.com

High-performance UPS systems
www.opti-ups.com