

Juner 21, 2006

RE: 16 Common Backup Problems and Mistakes

by Darren McBride

Summary: These are the most common back up problems and mistakes people make. It is the intention of this paper to give you advanced knowledge of these errors in hopes of preventing you from suffering the consequences of making these errors.

1. Tape errors, Tape Drive failures.

Probably the number one backup problem is simply caused by the unreliability of tape. Tape drives have a higher failure rate than most computer peripherals and the tapes have a limited life span. Many experts limit tape usage to 50 backups or 1 year (whichever comes first) due to tape wear. Failure to replace worn tapes will generate errors on most tape systems. Some types of tape drives (tape heads) need regular cleaning and dusty environments such as mines, construction sites, shops etc can eat tapes quickly. To avoid this use a high quality removable hard drive system, such as those found at www.High-Rely.com

2. Tape Logs are not checked.

All too often in a small business or office environment an automated backup is set up and the office personnel are told “It’s all working automatically at night – just change the tapes and everything will be OK”. This is the worst advice anyone could possibly give. From long experience we know that automated processes tend to fail over time. For this reason it is critical to review the “log” files after EVERY backup. Log files are available in most backup software and will allow the operator to see if files were skipped or other errors occurred.


3. Allowing users to leave workstations running critical software at night.

Most tape software will skip open files, which are often the most critical files in the enterprise. Many small business executives have fallen into the habit of leaving their computer on at night. While this practice is fine, leaving software up on the screen can wreak havoc on the nightly backup. If even one user consistently leaves an accounting application open, it is very possible that the most critical data in the enterprise will not be backed up (will be skipped). Reviewing the backup logs is critical to ensure this doesn’t happen.

4. Backups fail due to incorrect media or because previous job ejected tape.

Unless an expensive array of tape devices is used, tape media must be changed daily to reduce wear and ensure a wide, predictable variety of restore points are available. Some backup software (such as Symantec’s Backup Exec) allows a tape or media to be ejected when the backup is complete. It is common to see backups fail because the previous job ejected the tape and the operator failed to change it. Ejecting the tape helps the operator know that the job has run at a glance, but it can cause problems if the backups are unattended for several days. Too frequently the person who is tasked with changing the media doesn’t understand the importance and leaves without delegating another person to take over. If the tape isn’t ejected previous backups may be inadvertently overwritten. Another common failure with Backup Exec is that it may refuse to backup even if there IS a tape installed...but it is not the “correct” tape. Once a job hangs it can affect subsequent jobs and nightly

backups stop happening. Most of these issues are avoided with the High-Rely multi-bay devices, which provide the ability to backup to a different drive each night of the week. The product line and features can be reviewed at www.High-Rely.com.



5. Failure to buy or use agents for “always on” applications.

Backup software can be expensive. In small businesses the backup solution can cost as much as the fileserver itself! Many organizations are either not advised or refuse to purchase needed additional “agents” or client licenses that would allow them to backup more than one server and/or to backup always on software such as Microsoft Exchange or SQL. Even when “workaround” solutions are used, the organization may unwittingly cause themselves problems. For example, refusing to buy an Exchange agent will typically make it impossible to restore a single mailbox or message. Any restore that might be possible would require restoring the entire organization’s mail store. This means that if even one user loses an important mail that must be restored, all the other mail users would face the need to “roll back” their mail to that same point in the past, losing all of their more recent mail.

6. Relying on partial backups.

Over and over again IT professionals hear people say “I really don’t care about anything on my computer except for my _____ (fill in the blank – usually its accounting data, documents, photos.... or whatever). The truth is, those folks often don’t remember how many hours of work went into miscellaneous documents, spreadsheets, web page bookmarks, utility software and other things stored on the average computer, not to mention the cost of reinstalling the operating system and configuring it from scratch. In very small installations, not backing up the operating system may be acceptable on the theory that the OS can be reinstalled in a matter of hours. However, most businesses should backup

everything, including the operating system’s “System State”. It’s just too stressful and expensive during an outage to try to locate, reinstall, and reconfigure Windows, all the user accounts, and all the various software that has been added through the years.

7. Relying on incremental or differential backups with tape.

In a “differential backup” a typical scheme is to do a full backup once per week and on subsequent nights copy only those files which have changed since the last FULL backup. This means to restore you will require 2 media: the last full backup plus the last differential backup. An “incremental” backup is similar except only files changed since the last incremental backup are copied. Even more tapes or media will be needed. For example, when using an incremental scheme, if the last full backup was Monday and a failure occurs Friday morning you would need to restore the Monday tape and then the Tuesday, Wed, & Thursday tapes in order. If even one of these 4 tapes has an error, the possibility of an unsuccessful restore is quite high. Many IT professionals feel that although incremental & differential backups are much faster, inherently unreliable media such as tape makes using them a mistake. It is a less serious problem when more reliable media such as hard drive is used, especially if a more modern software package that is using a database structure to track file changes is also in play. However, it is always a good idea to fully verify to insure full recoverability.

8. Media or Data are not transported offsite on a regular basis.

Fire, flood, theft, employee malice etc can result in loss of critical data at work. Many organizations are under the mistaken impression that putting tapes in a fireproof vault at night is sufficient protection.

However, there are two types of fire safes on the market today; those for paper document storage (insurance papers, wills, checks, receipts, etc), and those for data storage or computer media. Data Safes carry a UL 125 or UL 72 Rating while document storage safes carry the UL 350 fire rating.

Remember – you can't store data tapes in a fire safe rated for paper because they will melt! The business should also think about regional disasters (such as Hurricane Katrina) that might affect both the business and a nearby residence that is being used to store offsite backups. A flood or serious earthquake may render large areas of a city uninhabitable and officials may refuse to allow admittance for extended periods of time. For these reasons it's important to delegate someone to periodically remove a backup media and bring back offsite media as needed and that the offsite location be as far away from the main site as is practical. An alternative is to setup Internet based backup to automatically transport critical data to another location.

9. Failure to have sufficient media for a good history.

A common failure scenario involves a database corruption in an accounting package that goes unnoticed until it's time to "close" the books at the end of a month. Only then is the problem discovered.... and each of the last 7 days of backup turn out to have the same error. Errors can creep in over time. Multiple copies (versions) of data should be retained so that if corruption occurs and goes unnoticed that older copies of the data can be accessed. Try to retain a minimum of 2 weeks of daily backups with additional monthly versions going back for at least 3 months.

10. Failure to setup removable hard drives correctly.

Software such as Symantec's Backup Exec requires hard drives to be created as virtual backup "devices". If the "device" in Backup Exec is chosen as a "Backup to disk" device instead of "Backup to Removable Disk" then changing media will result in hung jobs. Always choose the latter for removable hard drive devices.

11. Backing up the wrong data or drives.

Sometimes data is moved from one place to another by technicians who are solving problems such as hard drives filling up. Other times only specific folders are backed up due to limited space on the backup media and new applications are thus excluded from the backup. It's worthwhile to review the backup periodically to see if mission critical data is being missed.

12. Failure to consider the need to archive.

"Archiving" here refers to the need to mothball data over a long period of time. US federal laws such as HIPAA (for the health care industry), Graham-Leach-Bliley (for financial services industry), and Sarbanes-Oxley (for Public companies) require businesses to retain data, including email, for up to 7 years or more. In addition, defending a business against lawsuits can sometimes require a business to access old data. Whatever backup method is used, an attorney should be consulted to determine a reasonable length of time to retain business data.

13. Failure to enforce a policy that company data is stored on the server



In most companies, nightly backup is done only on the fileserver. However users will often save things in “My Documents” or other locations on their local hard drives. In addition to these documents not being available to other users in the business, the obvious implication is they don’t get backed up. Both technical (group policy) and administrative (business procedures) policies should be setup to discourage creation of data on local hard drives. As a minimum, users should be required to copy local documents to the server at logoff time.

14. The failure to backup desktops.

This common backup error is related to the one above about storing data on the server. According to IDC, up to 60% of all data is located on corporate laptops and desktops NOT on the fileserver as one would expect. The same study suggests that each year 10 percent of laptops are stolen and 15 percent fail. Yet most businesses backup only the fileserver and many companies take no steps to insist that data on “client” machines be copied to the server. The corporate backup strategy should ask users to leave desktops on at night. Licenses for the backup software should be purchased to allow backups of all these desktops to be done.

15. Trusting Mirroring or RAID as a backup.

It is a mistake to think that RAID 5 (Redundant Array of Inexpensive Disks) will protect data in the way a backup will. Although RAID will protect the business against a single hard drive failure, it does nothing to address multiple drive failures or a user who deletes something. Spyware, viruses, user error or maliciousness, the fact that there is a single copy of the data and it remains on-site, and a host of other problems mean that RAID is NOT backup and should never be considered one! As the graph below shows, fully 56% of data loss is not related to computer hardware failures at all. Finally, as pointed out earlier, you do not have a backup if your data is not offsite.



53% is “Soft” error (Human + Software + Virus)
47% is “Hard” error (Hardware + Nature)

Source: Ontreck Data International

Figure 1 – What are the common causes of data loss?

16. Changing the admin password without modifying the backup

This one surprises many people but it is actually the cause of a lot of IT service calls related to backup. Many backup software packages use the “security context” of a user to perform a backup. Whether correct or not, in small business the “user” chosen to perform the automated backup task is often the administrator. Due to various security concerns, a business will periodically change the administrator password and often will forget to do it on the backup software. This “breaks” the automatic nightly backup. The built in backup software in Windows (NTBackup), as well as Symantec’s BackupExec and many others require you to change the password inside of the backup software any time the backup user’s password is changed.



Appendix A: Characteristics of an Ideal Backup

Conclusion

Like the gentleman who is welding his gas tank with minimal security precautions, running a business without a good backup strategy is inherently dangerous. We've discussed 16 common backup mistakes. The title of this paper could easily have been 17 mistakes because there is another "backup mistake" that almost everyone makes and that is this: **Most people fail to do a full test restore of their backup.** This is perhaps understandable because it is quite difficult & time consuming to do. Most people are reluctant to break what doesn't need fixing. Certainly it's true that doing a test restore on a production system can potentially damage the working system. Solutions such as booting to alternate (dummy) hard drives and setting up "virtual" servers for testing are techniques backup professionals use to overcome these issues. Highly Reliable Systems manufactures a complete line of removable hard drive devices targeted specifically at the backup market. Although many of the problems we've discussed are software and policy specific, the products and backup experts at Highly Reliable Systems can help you to overcome these and many other common backup pitfalls. Please visit www.High-Rely.com to review our product line. Specific questions about this paper can be sent to Darren@High-Rely.com

1. Backup should work reliably with minimal user intervention.
2. Backup should not impact your day to day productivity (ie slow the server down during working hours)
3. Backup should be as fast as possible and happen automatically.
4. The Backup Device should be inexpensive.
5. The Backup Media (tapes, hard drives, etc) should be as inexpensive per Gigabyte per year as is possible
6. Data should be available to restore quickly using "random access" to any file. Ideally it should be so easy end users can do it.
7. Multiple copies (versions) of data should be retained so that if corruption occurs and goes unnoticed that older copies of the data can be accessed.
8. Multiple copies of data should be available on multiple different media so a single or double failure will not result in data loss.
9. Backup Media or data must be transported off site. If data is not taken off-site consistently it's not a real backup.
10. The entire network backup should fit onto one media or media set that can be accessed without human intervention so the backup operation can be done at night and/or unattended.

11. Any good backup should be capable of dealing with open databases and operating system files. Examples of “problem” areas for some backup software include active directory (On Windows domain controllers), Microsoft Exchange, always on SQL databases, Sharepoint and similar products.
12. Some backups should allow “granular” restore for organizations that require them. For example, many backups will backup the entire Exchange email database (the “data store”) but are unable to restore a single mailbox or a single user’s email. Instead the restore is “all or nothing” in which everyone’s mail in the entire organization can be “rolled back” to the specific time the backup was made. Similar granularity may be needed to restore individual files in products such as Microsoft Sharepoint.
13. In secure environments the backup should retain security settings (NTFS permissions) on files and folders, so that after a restore, the security is exactly as it was.
14. Some companies will want the Backup solution should have a simple “disaster recovery” option to speed up recovery. This option allows for a “bare metal” restore in which a server may be booted from media such as a CD and the data restored without the additional (and typical) step of installing the Windows operating system first.