Back Up Your Data or Backs to the Wall?

R.D. McDowall, McDowall Consulting, Bromley, Kent, UK.

When was the last time you backed up your computer? Thought so…

The road to hell is paved with good intentions — I should have backed up my data, but I didn't. I want to focus in this column on an essential part of the process of ensuring data security and integrity: back-up and recovery. This topic deals with the first stages of your disaster recovery process: the use of tape storage for data files, and system and application software. This is not to be confused with archive and restore that deal with the long-term storage of data on more permanent media such as CD-ROM. The organization of data is usually different in the two instances: back-up deals with all data on a disk whereas archive covers data organized around specific work packages.

How Important are Your Data?

The importance of the data stored on the computerized system or network determines how often your back-up and recovery procedure is performed.

How critical are the data? For critical data, the intervals between back-ups and the type of back-ups performed will be greater than low priority systems where back-up may be made at a lower frequency.

How often do the data change or are new data acquired? Data systems that acquire new files regularly (e.g., chromatography data systems) or manipulate data extensively will need more frequent back-up than systems in which change is much less. However the spreadsheets used instead of laboratory information management systems (LIMS) or chromatography data systems (CDS) capabilities are just as important to back up.

What speed of recovery is required? Can the system be restored within a working day with little impact or does it need to be restored within four hours? This will affect both the frequency and nature of the back-up and recovery schemes as well as the linkage with database transaction logging.

All of these issues need to be considered when designing the back-up and recovery process.

Regulatory Requirements Driving Back-up

The impact of electronic record and electronic signature regulations (1) also means that data must be backed up effectively to avoid data loss, as 21 CFR Part 11 has specific requirements that involve back-up and recovery of your chromatographic data.

- §11.10(b): The ability to generate accurate and complete copies of records in both human readable and electronic form suitable for inspection, review and copying by the agency.
- §11.10(c): Protection of records to enable their accurate and ready retrieval throughout the records retention period.

Thus back-up and recovery are very important functions that need to be validated to demonstrate that the overall process works and continues to do so after upgrades of software.

Furthermore, observations have been made during inspections, such as:

- Failure to comply with network system back-up procedures in that not all required back-up procedures were documented as scheduled, showing lack of documented evidence that tape replacements were done.

or

- There were no daily incremental back-ups; back-ups were performed every two weeks. If the system fails, data acquired between back-ups will be lost. The firm did not have contingency plans in the event of a system failure.

What are Back-up and Recovery?

Back-up and recovery are focused on storing and restoring system, application and user files.

- Storing means copying from a source device, such as a disk on your CDS or network drive to a target medium (usually a magnetic tape).
- Restoring means copying from media containing stored files to the primary location of the files: usually a high-speed disk.

Roles and Responsibilities

There are two main roles involved with back-up and recovery in a client/server CDS environment in most organizations.

- end-users
- IT department personnel.

End-users are responsible for back-up and recovery, as it is their system and their data; these facts are often overlooked when a back-up schedule is developed. Responsibility can often be abrogated and a default schedule given that is not appropriate to the system or the data held on it. Users must be aware of their responsibilities in this area.

The IT department will usually perform the back-up and recovery of data. The schedule for the back-ups will be worked out in discussions with the end-users and this should be (but rarely is) recorded in a service level agreement (SLA) that outlines the roles and responsibilities of all parties and the schedule. In addition, a standard operating procedure (SOP) will outline the schedule of back-ups, electronic records generated during the back-up process and any recoveries performed. These must be maintained for a validated application.

Hardware to Help Data Security and Integrity

In many instances, data and application files are stored on a single computer disk; this means that there is a single point of failure, which could mean the loss of your CDS files. As a result, you should also consider the use of hardware options to consider improving data integrity and fault tolerance with the system. Usually, these hardware options are grouped under the acronym of RAID (Redundant Array of Inexpensive Disks). There are three options commonly available to computers and servers commonly used to hold regulatory data.

RAID Level 0: *Data striping*: This involves two separate drives where any data written to disk are broken into data blocks called stripes. These are written in sequence to both drives. The advantage of this configuration is speed but the disadvantage is that if a hard disk fails the stripe set will be lost and the data will have to be restored from back-up tape. Apart from the speed gains, there are few other advantages of RAID 0 and, therefore, for data security and integrity, one of the other two options should be selected.

RAID Level 1: *Disk mirroring*: Data are written to two drives that are configured identically. The difference between RAID 0 and 1 is that when a RAID 1 drive fails, the other drive contains an exact copy of the data and can be used immediately (See Figure 1). However, to replace the defective drive the computer must be shut down for the defective disk to be replaced. There is a single point of failure because there is usually a single disk controller for the two disk drives; if two controllers are used then this is termed disk duplexing. One disadvantage of RAID 1 (and this is exacerbated in RAID 5) is that 10 Mb of data requires 20 Mb of disk space. Nevertheless, given the cost of disk drives versus the value of the data stored on them, this is usually a minor problem.

RAID Level 5: *Fault tolerance*: achieved by disk striping with parity, is essentially an extension of RAID 1. If a single drive fails then the data on it can be recovered from the other two by using the parity checksum information. However, if two disks fail then it's over to the tape back-ups to recover the data. To implement RAID 5, usually three identical drives are needed and the operating system (OS) is set up to manage the set. In the normal operation of the computer, data will be copied across the disks with the parity checksums. If a disk fails, some vendors offer a hot swap option where the old disk can be replaced with a new (but empty) disk and the data on the failed disk reconstructed using the parity checksums on the remaining two disks.

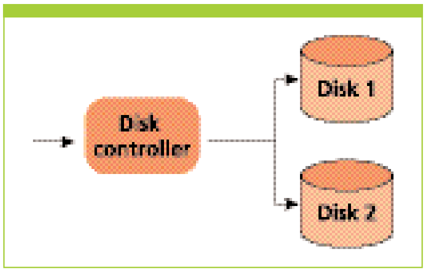

Figure 1: *RAID 1 disk configuration.*

Normally RAID 1 or 5 would be used to store data. However, if your data are very critical, then you must use RAID 5 with the fault tolerance aspects to reduce data loss. If the computer application needs to operate with greater than 99% of uptime, then you need to consider additional hardware features, such as dual processors and uninterruptible power supplies (UPS).

Options to Consider for Back-up

There are several options that can be used in developing our back-up strategy in addition to using hardware to mitigate the effects of a disk crash.

Full back-up: This involves a regular back-up of the system and is a complete copy of all system, application and user files on all disks to tape.

Incremental back-up: This is a regular, but partial copy of system, application and user files, identified by back-up profile, to tape.

Differential back-up: This is a regular, but partial copy of files that have changed since the last normal back-up.

Special back-up: This is a specifically requested copy of explicitly identified files to tape.

Most readers will be aware of the nature of a full back-up of the system but will be less aware of what the differences are between incremental and differential back-ups. Let's look at a typical example of a back-up.

- A full system back-up will be performed on the Friday evening or over the weekend when there are no users on the system. This will include all data and can also include the application and the operating system, although back-ups of the last two are usually performed on a less frequent schedule, as there is less change. The OS and application software can be physically separated from the data using separate disk drives.
- During the week, incremental back-ups will be made on Monday, Tuesday, Wednesday and Thursday.
- Each incremental back-up will contain the files that have changed since the last incremental or full back-up.

Assume a system disk fails on Friday afternoon before the next full back-up is scheduled. To recover the disk, the last full back-up is recovered and then the successive incremental tapes are recovered. Thus to recover the disk back to Thursday evening, the full back-up and four incremental back-ups must be installed. A failure in one of the early incremental tapes will result in lost data even if the later back-up tapes are perfect, amplifying the impact of any data loss, as I have discovered to my cost.

The differential back-up, in contrast, contains all changes since the last full back-up on the Friday night. Thus the Thursday tape will contain all changes in files since Friday's full back-up. After disk failure it only requires the full back-up and the latest differential tape to bring the system back to Thursday evening. The differential tape back-up will grow in size over a week, in contrast to an incremental tape that is relatively small in size.

The greatest advantage is that only two tapes are required to recover data from a differential back-up as opposed to a maximum of five with an incremental back-up.

Fundamental Back-up Activities

The first part of the process is to schedule back-ups for the CDS system. How vital are the data stored on the computer? Only the users can tell you! If you work in IT, don't anticipate what the users may think or say; get them to commit to the back-up schedule themselves. The user department will usually be paying as is the case in most pharmaceutical companies today. The cost of IT services is recovered by cross-charging for services provided to an individual system.

However as protection of electronic records is mandated by 21 CFR Part 11, then if the CDS system is GXP (good laboratory practice (GLP) and good manufacturing practice (GMP)) classified, the data must be held for the records retention period as specified in the applicable predicate rule. For GMP, this is the expiry date of the batch plus one year. However, many organizations retain CDS data for longer than this and to protect against litigation many are intending to store data indefinitely — meaning an upper limit has not been set. Criticality of the data will dictate the type of back-up and the frequency so that any potential data loss is minimized.

Hot or Cold Back-ups?

There are two basic approaches to back-up of computerized systems.

- Hot or on-line back-up takes place whilst the system is still operating.
- Cold or off-line back-up occurs when the system has been stopped and users have been logged off the system.

The cold back-up is generally thought of as the safest type of back-up, as the hot back-up requires the system to be buffered while the back-up occurs and the system updated when complete. The option you select is up to you and depends on the use of the system and the value of the data.

For instance, if you have a data system that is required to be available 24 hours per day, 7 days per week then a hot back-up of the system would be required. Alternatively, if the system was only required to be available 95% of the time, then a cold or off-line back-up approach could be devised where the system would be backed up when there were no users on the system.

Cold Back-ups

A cold back-up is scheduled out of working hours (e.g., say at 2 am) and is performed automatically, with the logs of the activity reviewed the next morning to confirm that all has gone well. The process is as follows:

Remove users from system: If you need to back-up the system during normal working hours warn the users of impending downtime, then ensure that all users are logged off and have their files closed. The system manager can disable further user access and also make sure the appropriate processes for back-up are running.

For normal hours or out of hours back-up the common process is:

Copy files to media: Using a software tool provided explicitly for the purpose of back-ups (either with the operating system or purchased specifically for the purpose), files are copied from their primary location on disk to back-up media, usually tape. Typical back-up applications are Backup-Exec and ArcServe.

Verify readability of backed-up files: You are placing your trust and your data in the hands of a magnetic medium; therefore, the quality of the back-up is essential. Verification confirms that the files have been copied to the tape correctly and that the tape can be read again. Verifying readability of files backed-up can include a comparison of the files on the back-up medium with the originals on disk and this gives the best confirmation. However, this can take time. Verification is important, as back-up is a vital component of data security and integrity. For off-line back-ups this step usually occurs before users are allowed access to their applications following the back-up.

Allow user access to system: If the back-up is verified as readable, the system can be returned to normal operations.

What happens if a back-up fails? Unsuccessful back-ups must either be rescheduled for the next back-up window or a back-up must be performed as soon as the problem is known. However, a cold back-up requires that users must be logged off the system and this may result in downtime. The system owner has to balance the potential loss of data against system downtime. In this instance it may be worthwhile considering a hot back-up schedule.

Hot Back-ups

Hot back-ups require a fast tape system to transfer the data to tape while the system is operational. There may be a slight but noticeable degradation of performance, but the availability of the system overrides this issue.

One way to overcome this is to have a second disk that is empty and the same size as the data disk on the server. Transfer the data from the operational disk to the empty disk; this is a relatively quick operation as the original disk reading, transfer via the internal bus and disk writing is far quicker than the disk to tape transfer. When complete, the image of the data on the second disk is backed up as if it were off-line. When the back-up is complete and has been verified, the data on the second disk can be deleted. The disadvantage is the cost of the additional disk and any associated service costs from the IT department. However, in my view, the benefits of this approach outweigh the disadvantages.

Media Management

Media are typically either DAT (Digital Audio Tape) or DLT (Digital Linear Tape) tapes. The former is cheaper and slower compared with the latter and if speed is required, then DLT is the current choice until something better comes along.

Media management is defined as the activities necessary to ensure that back-ups and restores have reliable media, where they need them and when they need them. This can take a number of forms, such as:

Media identification: Ensuring that all tapes are uniquely identified with a number, colour and, if an automated robot is used, a barcode.

Media rotation: Regular cycling of media used for back-ups includes replenishing supply and disposing of unreliable media (media are considered unreliable when they have been used beyond their normal supported life, when they have been found to be unreadable, when they have flaws making them unusable or an error is reported during back-up). This is key: do not think you can save money by reusing suspect media because you will pay a much higher price in the long run through data loss.

Logical media library: A catalogue of the back-up media with retrieval index, contents and location for each system.

Media audit: Verification that media can be found at the location specified, are readable and contain the data specified and are listed in the logical media library. Audits can be scheduled at regular frequencies to confirm that there will be no problem locating the appropriate tapes.

Dual locations: Once every two weeks to a month, full back-up duplicate tapes are made and they are stored in a separate location either on the site or off-site as a disaster recovery measure.

Manage back-up media generations: Depending on retention policies, determine which back-up generations can be reused and which must be saved. For example, if retention indicates that three months of the first full back-up of each month are to be saved, and it is the middle of September, the media from June can be reused.

Determine additional needs for new media: If media use is increasing because of greater volumes of data being backed up, more frequent back-ups, or other changes in the back-up profile, there may be a need for additional media. Plan proactively for this rather than run out of tapes and have no cover.

Restoring Data from Tape

Despite all the efforts of designing fault-tolerant hardware, there will be a time during the operation of any system that anything from a single file to a whole disk will need to be restored. This is where the appropriate tape is invaluable, assuming the back-up has been performed correctly and the tape can be read.

Request media from library: You will need to identify the tape or tapes that the data are on and bring them to the tape unit for the system. A restore request will usually indicate the file(s) to be restored and the date of the back-up. Thus, the media request resulting from this process will identify the media to be used

Execute and verify restore: Using the correct tape, identified through your super-effective library catalogue that you validated before it became operational, the file or data requested are restored to your system. Of course, we do not forget to verify that the recovery has worked. Database recoveries can be a little more complex than simple files. For example, a database recovery might entail recovering the log files (e.g., redo logs) following a restore.
Return media: Tape(s) are returned to the library.

Summary

We've looked at the rationale and process for back-up and recovery as well as some steps using hardware that you can take to make your computer system more resilient. However, we have not considered the records generated by back-up and recovery and the written procedures required under GXP regulations.

References

(1) 21 CFR Part 11, *Electronic Records, Electronic Signatures Final Rule*, Federal Register, (1997).

Bob McDowall is Visiting Research Fellow in the Department of Chemistry at the University of Surrey, and Principal of McDowall Consulting, Bromley, Kent, UK. He is also a member of the Editorial Advisory Board of LC•GC Europe.