



The Cotton Wool around data backups

A Management White Paper

Operational requirements

Infrastructure requirements

Infrastructure costs

Types of data

Types of Server

What should be backed up?

By John M Hudson

Director of T360

February 2003

Operational requirements

An application provides a functionally based service to its users; this aspect is usually clearly defined, providing the business case and ROI (return on investment) in order for a project to proceed. This service, however, has a further operational requirement; disaster recovery, often considered at the end of the deployment phase and often only addressed if funds allow. Deterministic recovery is not merely the domain of the large corporate operating 24 x 7 with remote warm standby sites. It is appropriate for all businesses with a level of dependency on IT – in short all businesses. Recovering from a disaster requires the reinstallation of the operating system, application, configuration and data onto a server within a given time frame which has been agreed with the business and deemed appropriate to meet operational objectives; a back to service agreement (BSA).

The BSA might be measured in minutes, hours or days and differ across applications and services. Even a corner shop, whose accounts PC can be down until the end of the month without impacting on the business, needs applications, configuration and data to recover from a disaster. It is, therefore, the premise of this document that the ability to meet the BSA is the prime objective of the backup process.

Infrastructure requirements

Often deemed an overhead, infrastructure does not provide additional functionality to the users and as such is a cost that reduces the ROI. LAN bandwidth provides access to services for users. This is a necessary part of any solution, but a substantial hike in this “merely” for backup purposes should be avoided. Minimising these costs is key to the viability of the project. This is only one example; tape libraries, storage and WAN bandwidth are all costs to be absorbed.

These stresses are increased by the simplistic philosophy that in order to remove all risks, a “backup everything” policy is adopted. It is felt that a serious review of this approach is required, whilst balancing this with the operational requirements. In order to address these issues, it might be necessary to adopt system architectural changes and mandate additions to the change management procedures.

The purpose of this document, having already established the operational requirements, is to examine the infrastructure required to support the backup and restore processes, along with the costs, savings and benefits associated. In short, how can we meet our objectives, improve the service to users and reduce costs?

Infrastructure costs

A file system consumes storage loosely divided into a number of areas:

- Operating system and swap
- Application software
- Temporary files
- Configuration files
- Data

In the classic “backup everything” approach, this can amount to a substantial amount of data. The question is how much of this is essential when meeting the real objective of server restoration? The more that is backed up, the greater the cost:

- Bigger tape library
- More LAN bandwidth
- More time to backup and thereby restore – poorer SLA

The Cotton Wool around data backups

Some of these issues can be addressed by adding more tape libraries to offset LAN bandwidth, however, this results in:

- Cost of more tape libraries
- Increased software license costs
- Additional installation
- Greater maintenance cost
- More support staff

Moreover, if the data cannot be delivered quickly enough, and tape drives are now 10 times the performance of but 5 years ago, the tape will operate in start/stop mode, and not stream. This will reduce its reliability.

This is the common evolution of a backup solution, as systems get larger, aided and abetted by tape and software vendors alike. An alternative approach is to upgrade the architecture to say a SAN (storage area network), but again there will still be:

- Some increase in license costs
- Fibre, interface cards and hubs
- Additional installation

Whichever way this is viewed, it results in an increase in costs. The key, therefore, is to analyse the contents of the file system to establish which elements are essential to meet the goal of system restoration.

Types of data

The manufacturer of both the operating system and application software come on CD, which is suitable for installation. There is no benefit in backing up these elements of the system. T360 has given further consideration to this aspect of system restoration and recommends the use of such products as Sun's 'Jumpstart' or 'Ghost' for Windows, to restore this software onto a virgin system. This approach provides a rapid and deterministic mechanism for bringing the underlying software of the server back on line. It should be noted that this process does constitute an element of the time encapsulated within the BSA.

Swap space, by its very nature, is transitory and hence does not need to be backed up. The same is true of temporary files, though in certain cases this needs further analysis. Temporary files are often created whilst running, say a batch process and, thus, if this task is incomplete at the time of failure it must be restarted. The resulting temporary files are, therefore, of little consequence.

Configuration files only change when a new and tested system is rolled out; there is, therefore, no benefit in backing these files up within the standard schedule of data. The most effective use for restoration in this particular instance is to take a 'baseline' backup of these files on completion of any upgrade or change. It is recommended that at least two, and preferably three, copies of this baseline configuration be archived in order to protect against flawed configurations or corruption during the change management process. This provides version roll back. The amount of data associated with configuration files is often modest and as it can be scheduled outside the standard backup window, normally as a manual process, it has little impact upon either LAN or tape library requirements, if this method is adopted. It should be noted that change management procedures must incorporate this baseline backup and it is imperative that these are adhered to.

Data can be sub-divided into five distinct groups:

1. Key operational data; raw data is the only information that it is imperative for the restoration of a given service and, hence, must be backed up.

The Cotton Wool around data backups

2. Historical data, by its very nature, does not change and once written can be classed as read only. If this data is housed in a database then an incremental backup will only backup data added since the previous backup. The only time this process fails is when historical data is appended to a file. Unless this file can be partitioned by day or week, it must be considered a potential design flaw in the software architecture. This must be managed to minimise the impact on backup and restore operations. So, whilst the historical database might be large, say a year's data, only 1/365th is actually backed up each night. The downside of this approach, if taken to its logical conclusion, is that to restore the data the initial full backup followed by 364 incremental backups, must be applied to restore the system which could be a very time consuming exercise. It may be more appropriate, therefore, to do weekly or monthly full backups to reduce this process. It is also worth enquiring as to the BSA for the restoration of the historical data. Experience has shown that live systems, and consequently their associated data, require a rapid restoration process. However, systems containing historical or archive data can often be brought back in a more leisurely fashion. This effectively is turning a single application into two, each with its own BSA to meet their respective operational requirements.
3. Data archived to meet regulatory requirements. This probably does not need to be backed up on a regular basis and, therefore, imposes little burden on the tape library and infrastructure.
4. Duplicate files are commonplace, either as multiple copies of the same document, or as attachments to emails broadcast to colleagues as a circular. Software tools are available to identify these and, in the case of attachments, reduce it to a single 'pool' copy.
5. Much of the data housed within systems can be classed as transitory in addition to the temporary files discussed above. Much of the logged information is also of a transitory nature, it being invalidated within a modest period of time and, consequently of little business value. When analysing the BSA, there is no business benefit to backing up data that will be obsolete before or soon after the recovery has been completed. Data that can be re-constructed from the fundamental or raw data could also be classified as transitory. It might be acceptable, take compute time after recovery, to reconstruct multiple views by processing the raw data. The reduction in BSA and infrastructure costs providing the larger benefit.

It has been T360's experience that vast quantities of data can be removed from the backup schedule by the careful analysis of this type of data.

Having analysed the data to be backed up, the next imperative is to examine the backup schedule. In order to achieve a BSA, it may be impractical to use the technique of multiple incremental backups. Furthermore, consideration must be given to as how long the data must remain available for restoration before it is deleted and the tapes recycled.

Types of server

Whilst there are many and varied types of server, experience has shown that they can be crudely divided into three sets: operational, data full and hot standbys.

The operational systems can be viewed as black boxes and, therefore, only require baseline backups to manage their configuration files. These systems require only a modest investment in backup infrastructure.

Data full systems often have large databases associated with them and need to be backed up whilst live. To avoid the use of distributed tape libraries, with the inherent overhead and costs described above, it is recommended that storage area networks be utilised. Whilst this approach has its own intrinsic costs, if focused on only this type of server, often found clustered for further resilience, these additional costs can be shared. This approach facilitates the use of large shared tape libraries, can be used to focus on a single server to minimise the recovery time, and offers greater resilience.

The Cotton Wool around data backups

As detailed in T360's "Server Consolidation" white paper, a SAN with shared tape library, and consolidated storage, has further operational benefits. It is also worth noting that buying a single large tape library and storage system results in a reduced price per gigabyte than a smaller, more distributed, approach in addition to those accrued from support cost savings.

Hot standbys, sometimes on separate sites, raise different issues. Replication retains referential integrity but also enables the rapid duplication of corrupt data which, at a stroke, disables the fail-over strategy. It is at these times that carefully structured backup and restore strategies have an instant payback. Further, the potential benefit of being able to parallel a large restore, through multiple tape drives, over high bandwidth fibre, will be reduced from many hours to a fraction of that time.

Hot standbys or Host A and B scenarios have a further facet. On the failure of Host A, Host B takes over in a predetermined way, fulfilling the BSA. At this time the organisation is vulnerable. A further failure would result in a discontinuity in service to the users. It is, therefore, imperative to re-establish Host A as quickly as possible, within another BSA. Restoring data from tape at this time would be inappropriate, as the "live data" will have been updated in the interim period. Consequently, the architecture of the total solution must incorporate the ability to build Host A from Host B by reversing the data replication direction. The time taken to complete this process must be accounted for in the BSA, and the procedures to be adopted, clearly documented.

What should be backed up?

The results of the analysis are:

- The cotton wool 'backup-everything' approach is very wasteful; many servers do not need to be backed up as part of the standard schedule: for example, proxy servers, firewalls, and network elements.
- Only a proportion of the total amount of data needs to be considered for backup; the rest (O/S, applications, swap, transient data) does not need to be backed up.
- For each set of data, choose a method of backing up that is appropriate to the data and its restore requirements.
- Ask, "Is this data essential?" If not don't back it up.
- Carefully specifying the BSA for each set of data can further reduce the impact that backups have on the system.

This approach allows us to focus on the truly crucial business data and deliver a deterministic recovery process to meet operational objectives.

T360 believes that the results of this analysis, together with the resulting architecture, will provide a solution that:

- Is quicker
- More robust
- Improves the BSA and is thereby better for the user
- Require fewer support staff
- Generates savings from the scale of purchase

In short, this analysis has many long-term financial benefits.

The Cotton Wool around data backups

T360

T360 is an innovative solutions provider that assists organisations to transition IT from being a necessary component to a strategic collaborator within the business.

'Aligning IT to the business is a Process not a Product'

Based upon best of breed ITIL solutions, T360 enable customers to achieve an end to end visualisation of their business services delivered by IT through the areas of:-

- Business Service Management
- Business Activity Monitoring
- Customer Experience Management

T360's philosophy is simple; we listen, understand and deliver. Our customers trust us because we provide innovation, expertise and commitment.