Helping you piece IT together

Considerations for Network Backup

Copyright Notice
© 2005 BH Consulting IT Ltd. trading as BH Consulting, All rights reserved. This whitepaper is protected by copyright. Any reproduction of material of this whitepaper must be requested and authorised in writing from BH Consulting. Authorised reproduction of material must include all copyright and proprietary notices in the same form and manner as the original, and must not be modified in any way. Acknowledgement of the source of the material must also be included in all references. BH Consulting reserves the right to revoke such authorisation at any time, and any such use must be discontinued immediately upon notice from BH Consulting.

Disclaimer:
BH Consulting has made every reasonable effort to ensure that information contained within this document is accurate. However no representation or warranty, expressed or implied, is given by BH Consulting as to the accuracy or completeness of the contents of this document or any information supplied. Recommendations outlined in this document are based on accepted industry best practises.

The information contained in this document has been provided by BH Consulting for information purposes only. This information does not constitute legal, professional or commercial advice. While every care has been taken to ensure that the content is useful and accurate BH Consulting gives no guarantees, undertakings or warranties in this regard, and does not accept any legal liability or responsibility for the content or the accuracy of the information so provided, or, for any loss or damage caused arising directly or indirectly in connection with reliance on the use of such information. Any errors or omissions brought to the attention of BH Consulting will be corrected as soon as possible.

Any views, opinions and guidance set out in this document are provided for information purposes only, and do not purport to be legal and/or professional advice or a definitive interpretation of any law. Anyone contemplating action in respect of matters set out in this document should obtain advice from a suitably qualified professional adviser based on their unique requirements.

The information in this document may contain technical inaccuracies and typographical errors. The information in this document may be updated from time to time and may at times be out of date. BH Consulting accepts no responsibility for keeping the information in this document up to date or any liability whatsoever for any failure to do so.

BH Consulting are under no obligation to update this document or correct any inaccuracies or omissions in it which may exist or become apparent.

Table of Contents

1. Network Backup Issues

Many companies have implemented Local Area Networks to support their business. Due to the nature of modern, the information produced and stored on the network is critical to the commercial viability and survival of many companies. Not only is the cost of the data stored on the network invaluable, the cost of the number of man hours required to accumulate and manipulate the data into its present form is more expensive than the actual hardware upon which the data resides. Compounding the value of the information may be the proprietary intellectual value of the information developed within the company.

In the event of a disaster, the physical equipment employed by a company can be restored quickly, either by using spare equipment or by purchasing it from a vendor. However, the cost of retrieving data that has not been backed up properly or not at backed up at all, could result in a company at best losing the confidence of its customers as delivery of products are missed, or at worse, losing business.

The issues involved in backing up the information held on a network are often underestimated. This results in many people thinking that attaching a backup device to the server and backing up the data is sufficient. However, the backup issues facing many companies are not always clear cut. It takes a broad comprehension of the issues, from backup media options and methods of restoration, to select and implement a solid backup system. The issues facing companies when backing up critical information are as follows;

1.1 Backup Media

There are many different types of media that a network manager can use to backup data onto. The most popular media used to backup data is magnetic tape media. Advances in tape technology allows up to 250 GB of data to be stored onto tape, with some data compression systems this figure can go up as high as 1 TB GB. Other options available are;

- Removable disks
- CD/DVD Drives
- Optical disks
- Diskettes
- Internet Based Solutions

Each of the above technologies will be examined in turn.

1.1.1 Backup Disks

Backup disks allow for the speedy backing up and restoration of data. Backup disks can either be removable disks that are installed within the server, disks within a SAN (Storage Area Network) or disks within a Network Attached Storage (NAS). The data is backed up to these devices which provide for fast read and write operations.

Additional protection can be got by deploying traditional backup tape devices to back up the information held on the backup disks to tape. This allows for large amounts of data to be backed up from the backup disks at any stage without worrying about the impact the backup could have on overall system performance. The backup tapes can then be stored offsite to provide additional security of the data.

However the number of disks necessary to provide a proper backup cycle could prove prohibitively expensive.

1.1.2 CDs/DVDs

CDs/DVDs can now provide high storage capacities, up to 1 GB with some manufacturers. Because of their long life ratings, an estimated 10 to 25 years, optical disks are ideal media for storing archives or data such as legal documentation which needs to be stored for long periods of time. However, the issue with CDs/DVD based backups is the amount of CDs/DVDs required to store the large amounts of backed up data.

Though ideal for perhaps backing up small workstations, this technology would be unsuitable for use with companies with large networks due to the potential large amounts of data that would need to be backed up.

1.1.3 Floppy Diskettes

Advances in certain diskette drives and diskette media now offer capacities of up to 40 MB per diskette. Though ideal for perhaps backing up small workstations, this technology would be unsuitable for use with servers with drives of multiple Gigabytes or indeed Tetrabytes.

1.1.4 Internet Based Backups

The advent of cheap high speed and high capacity Internet connections allows companies to backup data over the Internet to dedicated backup hosting centres. These centres use a central storage facility that enables clients to backup critical data. This enables the client to not only backup their data securely but also provides the extra security of having the backed up data stored offsite. The backed up data may also be encrypted whilst in storage to provide for its protection. Although a convenient way to back up data while also eliminating the need for offsite storage, there may be issues regarding the speed for restoring large amounts of data over an Internet connection. In particular, if the Internet connection is slow or has a high contention rate.

1.2 Tape Units

Magnetic tape units have proven to be the media of choice for most network managers as they offer high capacity backup solutions at a cost effective price. The most popular types of tape drives for network backups are Digital Audio Tape (DAT), 8-millimetre (mm) cassette, Digital Linear Tracking (DLT), Linear Tape Open (LTO) and Advanced Intelligent Tape (AIT) .

8mm Tape Drives

Backup systems that use 8mm tape are based on the Video8 cartridge format developed by Sony. The manufacturer, Exabyte, has licensed the technology from Sony and is currently the only manufacturer of 8mm data drives.

These drives use a method of recording called helical scan in which a rotating read-write head records data in short, diagonal tracks on the tape. The current capacity of 8mm drives range between 2.2 GB and 5 GB. However due to their low capacity, large sizes and slow speed 8mm tape drives are becoming less and less common.

DAT Tape Drives

Digital Audio Tapes are 4mm wide and come in a cartridge that is similar to the 8mm cartridge, but is much smaller. There are two competing standards for writing data to DAT drives: DataDAT and DDS (digital data storage). DataDAT is an update-in-place format that basically treats the tape as a random-access device such as a disk drive. DDS is a backup format promoted by Sony and Hewlett-Packard and has been adopted by a large number of tape drive manufacturers and integrators. DDS drives are manufactured by a number of vendors, including Archive Corp., WangDAT Inc., WangTek Inc., and Sony/HP. DDS 2, the latest version of the DDS standard, is emerging as the format of choice for writing data to DAT drives.

Like 8mm tape drives, DAT uses helical scan recording and currently provides up to 80 GB of capacity with 170-meter tapes. The DDS specification includes a high degree of error detection and correction. In addition, DDS supports a high-speed file find mode of up to 5MB/s that lets users find a file on tape in an average of 20 seconds.

DAT, especially the DDS format, provides the capacity for performing full backups on many of today's SME type file servers. DAT's ability to quickly find and restore files, multi-vendor support, and the availability of DAT changers makes it a very attractive choice for LAN backup.

Digital Linear Tape (DLT)

DLT Technology was developed by Digital Corporation (now Hewlett Packard). DLT drives use a linear recording method that places data in longitudinal tracks rather than the diagonal ones used in helical-scan drives. Data transfer rates approach 16 MBps with a 2-to-1 compression ratio, and reliability is often better than that of helical-scan technology drives. A full 25 percent of data on DLT drives is dedicated to error detection and correction. DLT drives are best suited to medium capacity backup requirements for where data is in the region of ten's of Gigabytes.

SDLT or Super DLT is an extension of the DLT technology and enables higher storage capacity.

LTO (Linear Tape Open)

The Linear Tape Open (LTO) standard is an open-format technology jointly developed by HP, IBM and Seagate. By developing an "open" standard the above companies provided a means that tapes and drives from different manufacturers are compatible with one another. Historically, tapes could often be read only by the drive that wrote them. Many companies faced issues when upgrading or changing their tape drives to find that they could not read or restore from old tapes used with the old tape drive unit.

LTO uses linear, multichannel, serpentine (back-and-forth) recording on 0.5-in. tape with a magnetic servo for error correction and hardware data compression. An embedded electronics module can store and retrieve usage and other information about a cartridge. LTO technology was originally announced in two variants, Accelis and Ultrium, aimed at speed and capacity, respectively. However, there was no demand for Accelis, and it has since been withdrawn.

The Ultrium format is a direct competitor to Super Digital Linear Tape (SDLT) that uses a single tape spool inside a cartridge. The current generation of Ultrium tapes can store 400GB of data in native mode, or 800GB if compression is used.

LTO Ultrium drives are becoming more and more popular especially in the high end capacity market and are often used in conjunction with disk based backup solutions for archiving and offsite backup. However, for low to midrange requirements they can remain an expensive option.

Advanced Intelligent Tape(AIT)

The Advanced Intelligent Tape (AIT) standard was developed by, and is exclusive to, Sony. AIT uses helical-scan recording on 8mm tape, similar to that used in Hi-8 video camcorders. With a higher bit density and narrower tape, AIT cassettes are smaller than other tape cartridges, allowing for tape libraries that hold more data but take up less space.

AIT cassettes include a memory chip inside the media cartridge to record and store format and file-location information. This lets AIT tapes load faster and cuts file search times in half.

AIT offers compression averaging 2.6:1 across multiple data types, vs. the 2:1 average for normal compression and can have up to 500GB native capacity, 1TB compressed. However, due to its proprietary nature and higher costs it is not a popular solution for companies with small to medium backup requirements.

1.3 Tape Systems

There are two primary classes of tape system vendors: manufacturers and tape drive integrators. Manufacturers such as Archive, WangTek, Exabyte, and WangDAT build the drive mechanisms. Tape drive integrators such as IBM, Mountain Network Solutions Inc., Maynard Electronics, HP, Tecmar, Dell and Palindrome Corp., combine those tape drives with host adapters, cables, software, and value-added services.

Although it is possible to buy drives and controllers separately, it is more advisable to buy a solution from a tape drive integrator. This will ensure that the interdependencies of the hardware components of the backup solution will work together. It is also very important to ensure that the backup software purchased has been tested and certified to run on the hardware chosen. Backup systems are the wrong components to save money on.

1.4 Capacity and Backup/Restore

With file server disk storage constantly increasing, backup system capacity is a major issue. Today, multi Gigabyte hard disks are common, and it is not unusual to find file servers with over 100 Gigabyte disk capacities. A file server should be backed up when all users are off the system to ensure no interruption to or degradation in services; in most cases, this means at night. Ideally, a backup system should have a large enough capacity to back up an entire file server, which allows for unattended backup.

1.4.1 Compression

Fortunately, tape capacities have kept up with server capacities. In addition, using data compression can increase the capacity of backup systems. Although compression ratios vary for different file types (and can go as high as 98 percent), on average, compression can double backup media capacity. Compression schemes, however, usually impact backup system speed negatively. If possible data compression should be implemented at the backup hardware rather than the software as hardware compression is in general faster than software compression.

1.4.2 Auto-Changers

Auto-changers for tape drives that hold five to 12 tape cartridges are available. The auto-changer will insert a tape into the backup unit, when that tape is full it will be removed and replaced by the next tape. Thus a 12 tape auto-changer can give the ability to backup up 12 by 400 GB tapes or in other words up to 4.6 TB. However at rated speeds of around 210 GB/hr, to backup this amount of data would require around 23 hours or nearly a day per backup. .

1.4.3 SCSI Chains

SCSI cards allow up to 14 devices to be attached per SCSI adapter, thus 14 units could be chained on one adapter card giving similar facilities to those provided by the auto-changers. With both the auto-changers and the chained SCSI devices it is important to note that these scenarios will only work with software that supports such solutions.

1.4.4 Time

With small local hard disks, excessive backup or restore time is usually more of an inconvenience than a major problem. However, with high-capacity file servers, being able to properly back up within a limited time period is critical, as is the ability to restore a single file or an entire file server in a timely fashion.

A backup system should be able to back up required files within the amount of time allotted. More importantly the ability to restore a file or files quickly from tape is a major issue. The time taken to restore a full tape back to disk will be at least the same amount of time as that to backup the tape. This needs to be factored in when deciding on the preferred solution.

1.5 Security

A major concern when backing up data is the security of the data once it has been backed up onto the backup media. This security is not only concerned with physical security of the media, but also with logical security.

(i) Physical Security

The importance of securing backup media in a safe place cannot be over stressed. If the backup media is not stored in a secure place away from the servers then it could also be destroyed in the event of a disaster. Backup media should ideally be stored off site in a proper storage facility designed to store computer media, i.e. the facility is kept at a proper temperature and free from magnetic fields and humidity. The off site storage facility should also be secure from intruders. There are a number of firms offering facilities such as described above and costs vary according to the amount of media to be stored and how regularly media need to be transported to and from the storage facility.

A cheaper solution is to store the backup media in an air tight fire proof safe. However it should be noted that, even though a safe may be rated as fire proof it may not protect the backup media in the event of fire. This is due to paper having a flash point at a temperature higher than that of backup media. It has not been unknown for fire proof safes to be opened after a fire, to discover that though the paper stored there was intact the backup media had melted into an unrecoverable state.

(ii) Logical Security Another security issue is that of access to the data held on backup media. In order to read data from an unprotected media, all that is required is similar software and hardware that backed the data onto the media in the first place. It is therefore essential that in order to guarantee the security of stored data, the software that backed up the data can protect it from unauthorised access. This can be done by numerous means, such as

Encryption - This is where the data is encrypted into an unreadable format as it is backed up. To recover the data successfully the encryption key entered at the initial back up operation is required to decrypt the data and make it readable. There are a number of encryption algorithms available, some proprietary to the software manufacturer, and others such as the AES, RSA and IDEA algorithms which are industry standards. However it should be noted that some encryption programs are not as robust as others and this should be taken into account when choosing a product. Other issues to be aware of with encryption are that encryption can slow down the backup process. Also, if the original encryption key is forgotten or entered incorrectly the data will not be recovered properly.

Password - Some packages allow the backup to be password protected thus allowing only access to the backup when the password is successfully entered. This method provides some means of protection while ensuring that the speed of the backup is not impacted. As with encryption, if the password is forgotten, it may not be possible to restore the data.

Another issue to address is the preservation of the network security attributes attached to files. This ensures that when a file is recovered from the backup then the only users capable of accessing the files are the original users.

In order to preserve the integrity of the data held on the backup media it is essential that both the physical and the logical security is properly addressed.

1.6 Reliability

To prevent data loss, you must have reliable backup systems and reliable, tested backup and restore procedures. Reliability means that the backup system provides error-free backups and error-free restores. This is a hardware, software, and procedural problem. The best backup software in the world won't solve problems caused by unreliable hardware. In addition, if backup procedures are inadequate, then critical data may not be properly backed up.

Backup hardware should not break down often. Although Mean Time Between Failure (MTBF) figures are an indication as to how long the device should operate without failing, they are usually an estimate by the manufacturer. In addition, MTBF figures are usually quoted at a rated duty cycle, or percentage of time that the unit will be in use which may not reflect "real world" situations.

Backup hardware should also provide extensive error correction and detection and be able to bypass or block out defective sections of media during the backup process. Backup media can be damaged between the time files are backed up and the time files are restored. If the media is damaged, the backup hardware should be able to read beyond the

point(s) of damage and continue data recovery, even if the data in the damaged areas is lost.

Although many backup devices provide high levels of error detection and correction and the ability to recover data beyond the point of a major media error, much of today's backup/restore software, primarily for tape, cannot take advantage of this capability. In addition, many backup software packages provide inadequate verification capabilities and don't compare the files on the tape to the files on the disk. Most often verification doubles the amount of time taken for a backup.

However, even the most reliable hardware and software will not make up for improper or inadequate backup and restore procedures. Backup procedures must be properly planned and should be documented. It is equally important to develop proper and effective restore procedures. In many organisations, restore procedures have never been tested, there are no written guidelines, and all too often the software required to restore files after a server hard disk crash cannot be located.

1.7 File System and NOS Support

There are many Network Operating Systems (NOS) available on the market. It is important when choosing a backup device and software that the Network Operating System supports the solution chosen. When choosing a backup solution, it is important that issues such as preservation of file attributes and file security is supported and that the backup or restore process does not place an undue load on the file server to the detriment of other services.

The ideal backup system should have the ability to properly back up and restore all files being stored on a file server, including system files and any extended file and directory attributes, file access dates, and file and directory rights information.

While most Network Operating Systems will support multiple clients with different operating systems, such as Microsoft Windows, Macintosh, Linux and UNIX, the backup software for a particular platform may not preserve the extended attributes of files from a different operating system than that on the server. A network incorporating servers with different operating systems may require a different back up solution for each operating system.

1.8 Cataloguing

Many traditional mainframe backup solutions offer extensive cataloguing facilities. However many network based backup solutions offer little, if no tape cataloguing. Most network solutions write a file to the header of the tape detailing what files are held on the tape. When a file needs to be restored this file has to be read into the memory of the server. This requires the network administrator to develop a method of locating tapes by label, date, or description, then locating files once a tape has been selected. Other backup solutions require the full tape to be read to find if the required file is on the tape. This can obviously add a lot of time taken to restore a file.

Ideally a backup solution should provide an on-line file history database that contains the archiving history for all files on a volume, allowing the ability to find and restore any backup version of a file.

1.9 Ease of Use

Another important issue when selecting a backup solution is to ensure that it is easy to use. With the explosion in the use of GUI interfaces such as Microsoft's Windows, LINUX, Apple Macintosh and many UNIX GUIs, a backup solution should ideally be GUI based allowing intuitive use. The ability to select files for backup or restore should be relatively easy. The type of backup whether it is a full backup, incremental or differential should also be simple to set up. The ability to secure backups by either using encryption or password protection should also be easy for network managers and/or operators to follow. Scheduling backups for unattended operation should also be intuitive and easy to set up.

1.10 Cost

The cost of a proper backup solution not only includes the cost of the backup hardware, but also the cost of the backup software, the backup media, the capability to store media off-site, the development of backup and restore procedures and the training of network administrators and staff in all aspects of the backup solution from procedures down to using the hardware. Very often the cost and nature of the data stored on a company's server is worth more than the hardware upon which it is stored, to reconstruct that data in the event of a disaster could be very costly, if not impossible, in terms of man hours and down time. When viewed against the cost of data loss in the event of a disaster, an effective backup/restore solution is very often trivial compared to that of losing all data.

By establishing proper backup procedures and using automated, unattended backup schemes, the labour costs for backup can be kept to a minimum. This is one area where purchasing high-capacity backup devices can produce immediate dividends.

1.11 Hierachial Storage Management (HSM) Systems

Several vendors also supply Hierarchical Storage Management (HSM) systems, which can integrate a network's entire storage system into a single, intelligently managed unit that encompasses individual drives, drive arrays, optical arrays, and tape jukeboxes. Generally, these systems can handle an unlimited amount of disk capacity, though practicality and performance on the Intel platform dictate that a new HSM server should be added for every 1 TB of storage.

HSM can be thought of as a multilevel fountain, in which data flows from one level to the next. The point at which it flows is determined by the age of the file; for example, if a document hasn't been read in a month, an HSM system can automatically move it from the on-line drive to a tape or optical jukebox, verifying a safe copy before deleting the file from the first drive. The HSM software keeps track of where everything is located, and fetches files when users or applications request them.

By setting different storage levels for different volumes and directories, network administrators can relieve themselves of constantly having to scan volumes and maintaining adequate free space. HSM systems are not cheap and are targeted at mid-sized to large network environments that can benefit from their extensive array of reporting capabilities, performance monitoring, and error tracking and alerting.

1.12 Server vs. Workstation Backup

There are two ways in which servers can be backed up, one way is to attach a backup device to the server (Server based backup), the other is to have the backup device attached to a workstation and to backup the server over the network onto the backup device attached to the workstation (Workstation based backup). There are advantages and disadvantages to both solutions

Server Based Backup

The main advantage of server based backup is that it is faster than workstation based backup for backing up and restoring files to disks locally attached to the server. Using a server based backup solution, unattended backups can be scheduled without requiring a user to be logged onto the server, thus improving security.

However, there are a number of disadvantages when using a server based backup. If there is a problem with the backup software this can have an adverse effect on other processes running on the server. If there is a problem with the backup hardware, it may require that the server be shut down and restarted to clear the problem thus affecting users on the network. Another problem with server-based backup is that, in the case of a file server crash, the backup software must be re-installed before files can be restored. Many server-based backup systems do not have a workstation-based equivalent to allow restores to a local drive or to a server running a different network operating system. This can create problems if data needs to be restored to a dissimilar server in an emergency or when migrating to another network operating system.

Workstation Based Backup

The main advantage of Workstation backup is that hardware or software problems do not affect file server processes or performance. Backup components can be removed and replaced without affecting the file server. In addition, files can easily be restored to other servers and local drives. In the case of a server crash, restore software can be run from a local drive, so it doesn't have to be loaded on the server before files can be restored. If a restore is done on a workstation local drive, the files restored are only accessible by the user of that workstation.

A serious limitation with workstation-based backup system is that it is generally slower than with server-based systems. Also, the backup station must be logged in to the file server in order to back up. Another consideration if choosing a workstation backup solution is the cost of purchasing a PC to use as the backup workstation.

1.13 Backing up Client Workstations

Workstations, PCs and/or laptops, can either store their data onto the file server or onto their local hard disks. If stored onto the file server then that data can be successfully backed up from the server. If the data is stored on the workstation it is then the responsibility of the user to back up that data. However, it is often the case that the user will not back up this data, which is subsequently lost if there is a problem with the workstation. Apart from requiring that all workstations are diskless workstations, therefore forcing the user to store data on the server, it is very difficult to manage data stored on users' local hard disks. There are numerous approaches that can be implemented to protect this data.

- One approach is to create a home directory for each user on the server. This will provide the user with an area on the server in which they can store their data and have it subsequently backed up. It is still up to the user to store their data onto the server.

- Another approach is to educate the user into backing up their data. This could be done onto diskettes using the CD/DVDs using the Windows backup command or onto a portable tape unit that is made available to users on request.

Both of the above approaches leave the responsibility for data, and very often sensitive data, with the user. If the user does not back up the files properly or as in most cases, if they do not back up their files at all, then that data is being exposed to the risk of being lost. In implementing the above solutions it is very important to stress and educate the user in their responsibilities for data management.

The other solution being offered by numerous backup software vendors is the ability to back up local hard disks onto a backup unit attached to a file server over the network. This requires an agent to be loaded onto the workstation to communicate with another process running on the server. These two programs allow the back up program on the server to backup files from the workstation to a backup device attached to the server. While this method allows for the centralised back up of all data on the network it does have a number of disadvantages that could make it impractical.

- One problem is security, in order to backup a user's files the user may have to leave their workstation turned on. As ideally the backup will be running unattended at night this will enable anyone to attempt to access the user's workstation. Also as the backup is being done across the network, any network analyser could collect the data as it is being passed to the server.
- The addition of an extra agent to the workstation can cause problems with incompatibilities with other programs and the using up of system memory.
- It could be as difficult to get users to use this procedure as it is for them to use the alternative procedures mentioned above.
- Backing up a large number of workstations over the network could increase the amount of network traffic and use up precious network bandwidth thus affecting other applications. Backing up remote workstation may prove difficult due to bandwidth congestion on the network, routers and firewalls.
- Finally to backup a large amount of workstations can be very time consuming and may require more than one tape, thus all workstations may not be able to be backed up at the same time. This could result in different workstations being backed up at different times.

1.14 Network Backup Procedures

It is important that once a backup system has been selected that both backup and restore procedures are put into place and that the necessary staff are made aware of how these procedures operate. The procedures should outline the roles and responsibilities individuals have in relation to backing up and restoring files. These roles and responsibilities should state who is responsible for starting backups, who is responsible for restoring files, when files can be backed up or restored and who can request restores.

Most importantly the procedures should state what to do when a disaster occurs. It is important to define what a disaster is and rate them in order of urgency. A disaster can be defined and rated from where an individual file has been lost or to where the computer room is off-line or has been destroyed. The procedures should be updated and tested regularly as all systems eventually change and staff turnover out of their areas.

It is important to document all the procedures so that no ambiguity remains for responsibilities and that procedures can still be followed if the original member of staff is unavailable. The documentation should contain;

- **Backup procedures**
  What type of backups are used and what the backup cycles are. Also if backup procedures aren't written down, they're forgotten.

- **Restore procedures**
  It is important that the documentation relating to restore procedures is clear and easy to follow, as file restoration at times may need to be performed by someone with little experience. It is also important that restore procedures are documented not only for single file restores but also for disaster recovery.

- **Test regularly**
  Like all other systems backup and restore systems can malfunction and it is important to test the systems regular basis. It is also important to test the restore procedures for different scenarios to ensure that different types of disasters can be recovered from as quickly as possible.

1.15 Backup Types

There are a number of different types of backups that the network manager can implement as part of the backup procedures. These are

A Normal or Full Backup

A full or normal backup is where all files stored on the server are backed up regardless whether or not they have been backed up previously. One of the advantages of normal backups is that they require the minimum number of backup sets and are therefore quick for restoring files, particularly if the restore is in the event of a disaster. Another advantage is that a full or normal backup is the most complete backup, i.e. all files are backed up.

The main disadvantage of a full/normal backup is that it can be time consuming especially on systems with large amounts of volatile data stored on them.

Incremental Backup	An incremental backup is where files that have been created, modified or deleted since the last normal or incremental backup are backed up. This can reduce the amount of time taken to back up. However the main disadvantages of this method are the larger number of backup sets required, and also, to recover from a disaster or in some cases just to recover a file, requires the restoration of the last full backup and every incremental backup since then.
Differential Backup	This type of backup will back up files that have changed since the last normal backup. This method is quicker than normal backups as fewer files are required to back up. However, if large amounts of data changes regularly this method can become nearly as time consuming as the normal backup. The main advantage is that for recovery only two backup sets are required, the last normal backup set and the last differential backup.
Archive Backup	An archive backup is where data that has not been used within a certain space of time is backed up onto the backup media and then deleted from the server. This allows for precious disk space to be freed up on the server. However it should be noted that archived backup media should only be used once and then stored in case the files need to be recovered at a later date.
Copy or Move backup	This type of backup copies selected files to the backup media but does not mark them as being backed up. This allows for the transfer of files from one server to another for the generation of a once off backup. This type of back up is particularly useful if tasks on the server, such as operating system or hardware upgrades, need to be carried out. This backup would allow the network manager to restore their system to the same state if a problem should arise.

Which backup method to use is dependant on issues such as the amount of data to be backed up, how often files need to be restored and what time is available for the backup to run.

1.16 Backup Rotation Cycles

How often a backup should be taken depends on a number of issues, but primarily, how often data changes and how valuable it is, are the main factors in deciding how often a backup should take place. Having more than one copy of a backup is recommended to guard against backup media failure. Whether to use full, incremental or differential backups is dependent on the factors outlined above.

It is common practise for organisations to implement a backup rotation cycle. This enables the organisation to re use the backup media at a regular basis while still protecting their data.

Backups could be scheduled daily and then recycled the following week. One of the more common rotation cycles to use is one that uses 20 backups over the course of one year.

A daily back-up of all data held on the server can be carried out every Monday, Tuesday, Wednesday and Thursday. This backup can be normal, incremental or differential depending on the needs of the organisation. The backup media for each of these days should be rotated weekly. The backup media should be labelled according to the day of the week.

A full or normal back-up which would back up all data and system files to be carried out each week on Friday. This procedure should be rotated on a four weekly basis. The backup media should be labelled as Week 1, Week 2 etc.

At the end of each month a full or normal back-up of all system files and data files should be carried out. This backup should be labelled with the month it was carried out and stored off site to ensure integrity. This will require two backups be carried out on the last week of each month, the normal weekly backup and the monthly backup.

Ideally a backup more than two days old should be stored off site to ensure it's integrity. The off site location should be a secure location with adequate safeguards to prevent damage to the backup media. Any backup being stored off site should at least be password protected and ideally encrypted.

2. Contact Us

Helping you piece IT together

If you wish to contact us or provide any feedback on this whitepaper you may do so using the following contact details.

Telephone : **+353-(0)1- 4404065**
Website : **http://www.bhconsulting.ie**
Email : **info@bhconsulting.ie**