



Microsoft

DATA PRESERVATION

*Mapping out a Reliable, Secure, and
Effective Backup, Recovery & Archival
(BURA) Plan for Small Business*

In the online-all-the-time world, business literally runs on data. Data – or more specifically information culled from volumes of stored files – is the resource that drives business innovation, expansion, and profits. Absent this information, business can and sometimes does come to an abrupt halt.

Yet in today's more virtual business environment where as many as 90% of all employees work remotely or in a branch office¹, files are in continual motion, passing between internal staff members and external partners, and through multiple client and server systems. With all this free flowing information, a major data loss is a real risk. Consider the fact that even with no precipitating external event, between 2% and 4% of all hard drives will fail. In some environmental conditions,



the failure rate approaches 13%². Factor in human error, natural disasters, and other incidents and the potential for a disruptive failure is even greater

Too many organizations have found out first hand just how costly data loss can be to their business. Research shows that businesses that lost data often suffered serious financial consequences, with nearly 33% reporting a precipitous sales drop and 20% admitting they lost customers³ Even a minor failure can become a big headache once you add up the expenses associated with an incident such as technical support required, lost productivity, and the value of the lost data itself. Some estimates peg just the average data loss cost at \$2900 per incident depending on the business value of the files impacted, and how the loss of access to that information affects operations and productivity⁴.

The onus is on the small business to take the most fundamental steps to protecting their information – copying data during a regularly scheduled back up so that in the event of a failure, an administrator can restore individual files. Yet, approximately half of all SMB with fewer than 250 employees are not backing up their servers every day, instead relying on weekly or

even less frequent file backups.ⁱⁱⁱ

The consequences of inadequate data protection can be devastating for a business of any size but small, often under-resourced companies pay a particularly dear price. Fifteen percent of companies with fewer than five employees that have lost data admit they were never able to recover the filesⁱⁱⁱ.

DATA IMPERATIVE

Unfortunately, this number is likely to be even higher for smaller businesses. A small business without an adequate backup solution or a practical plan to execute on quickly finds itself in an untenable position. Essentially, any company without a backup solution is gambling with its business data.

Small businesses that do not regularly back up their data cite a number of reasons for skipping this critical operation, including lack of expertise, uncertainty about which type of media to use, and concerns over the price of backup solutions themselves. Some small business owners also admit that the expectation that servers and other business resources be accessible 24 hours a day leaves an almost non-existent window during which they can back up their data.



Yet even businesses that do backup their servers often stumble when it comes to their approach. Should they use tape or hard disk? Should policies be based upon incremental and/or full backups? Which is the best media rotation strategy for their specific IT ecosystem and business requirements?

MEDIA MATTERS

With respect to the backup media, the most popular types, magnetic tape and hard disk each have their own distinctive strengths and downsides. Tape is a stable and reliable medium with a long history in backup. Traditionally less expensive than disk though that price comparison is beginning to shift as Disk prices fall, tape is also more compact and lightweight, making it easier to store.

Hard disks were traditionally more expensive than tape but that price point has been falling. Disks can be connected both locally through interfaces such as SCSI, USB, FireWire and eSata and over longer distances via transmission protocols such as Fibre Channel. Disks sometimes incorporate data deduplication, a method that gets rid of unnecessary data to free up more capacity.

Disk backups can run in the background with minimal impact on application performance. IT administrators can also set up very granular backups, choosing just to copy a select group of files, to minimize the impact on their server.

Hard disks are more fragile than tape, which becomes an issue when they are moved offsite. There are also questions about whether hard disks might degrade over time.

INCREMENTAL IMPROVEMENTS

All businesses should start with a full backup of all of their server files to serve as a reference copy. Companies can also choose to take one of several approaches beyond that to preserve just the changes that have occurred since the last full backup. Incremental and differential backups are perhaps the most straightforward and efficient methods. In the incremental method, just the files that have changed since the last daily backup are backed up to tape or disk. Differential backups record files that have changed the last full backup, excluding those that have already been written to media. Therefore, if the last full backup was on Sunday and it is now Monday, just the changes since Sunday are

| Data | Original System | Day 1 | Day 2 | Day 3 | Day 4 |
|---------------------|--------------------|-------|-------|-------|-------|
| Full Backup | Backup of Original | | | | |
| Differential Backup | | Day 1 | | | |
| | | Day 1 | Day 2 | | |
| | | Day 1 | Day 2 | Day 3 | |
| | | Day 1 | Day 2 | Day 3 | Day 4 |
| Incremental Backup | | Day 1 | | | |
| | | | Day 2 | | |
| | | | | Day 3 | |
| | | | | | Day 4 |

recorded. On Tuesday, the files that have changed since Sunday or backed up. In a differential backup, just the changes since the last backup are recorded, so on Monday just the changes since Sunday are backed up to media, and on Tuesday just the changes that were backed up since Monday are backed up, and so on and so forth.

Differential and incremental backups use less space on tape. They also allow the IT administrator to keep copies of different versions of the same file on different tape sets.

While incremental and differential backups are quick, there are a number of

downsides to using this method, not the least of which is that if any one of the tapes is damaged, full server restoration will be impossible. Another issue with incremental and differential backups is that to restore the server, an IT administrator needs both the tape with the most recent changes and the original backup tape.

A full backup, also known as a reference backup, copies all files from the server's hard drive. By capturing all the files on the server, the full backup provides the most complete source for a restoration in the event of a failure. This also gives the IT administrator a complete set of all the changes.

| Week | Mon | Tue | Wed | Thu | Fri | Satu | Sun |
|------|-------|-------|-------|-------|-------|-------|-------------|
| 1 | Son 1 | Son 2 | Son 3 | Son 4 | Son 5 | Son 6 | Father 1 |
| 2 | Son 1 | Son 2 | Son 3 | Son 4 | Son 5 | Son 6 | Father 2 |
| 3 | Son 1 | Son 2 | Son 3 | Son 4 | Son 5 | Son 6 | Father 3 |
| 4 | Son 1 | Son 2 | Son 3 | Son 4 | Son 5 | Son 6 | Grandfather |

One of the downsides to a full backup is that it is a much slower process than the incremental backup because the entire hard drive is copied. A full backup also takes up more space on the tape.

IN ROTATION

To maximize the use of their tape or disks while adequately backing up their server resources, many companies employ a rotation strategy that allows them to mix methods by “scheduling” how each individual piece of media is used. For example, as part of a rotation strategy an IT administrator can choose what kind of back up (full, incremental, etc.) as well how frequently it will be performed (daily, weekly, monthly) and when the tape should be removed to off-site storage for archival/disaster recovery purposes.

There are a number of common schemes for how to rotate media in and out. One of the best known and perhaps most pragmatic and effective is the Grandfather-Father-Son or GFS rotation strategy. GFS mixes daily incremental backups with occasional (weekly and monthly) full backups. Though initially designed for tape, GFS can also be used with disk. A company needs eight of whatever media they choose to start – four for incremental daily backups, three weekly and one monthly.

The way it works is one set of media (either tape or disk) is selected for the four business days, for our purposes, the first four days of the week (Monday through Thursday). These tapes, known as the “son”, are used for incremental daily backups. At week’s end, one son tape graduates



to father status. In subsequent weeks, the daily son backups overwrite the earlier copies on the same day. So unless the tape use on Tuesday has “graduated” it will back up files on the subsequent Tuesday.

A second set of tapes – the father – is set up for weekly backups that run on Fridays. Full weekly backups occur on the day of the week that the “son” is not copying data (e.g. Friday). These tapes are reused each week. One “father” graduates to being a “grandfather” once a month. The cycle actually kicks off initially with the “father” doing a full backup on the first Friday of the month.

The last set of media, the Grandfather, supports a monthly backup for three months. The Grandfather tape does a full backup on the final business day of the month, and then is retired to an offsite location.

There are other rotation strategies based on mathematical models (The Towers of Hanoi) and other highly complex schemes such as incremental media. However, the GFS rotation cycle is popular because it is relatively easy to implement, maximizes utilization of the media, and provides excellent disaster recovery capabilities in the event of an outage. The only real concern with this scheme is that in

its quest for efficiency (where all but the grandfather tapes are reused multiple times) there is the danger that the tapes may wear out.

Stable Foundation OR Data Recovery: Why does it matter?

Whatever the media, method, and rotation schedule a small business uses, it is critical the company be diligent in sticking to the schedule and maintaining their media adequately, including moving the appropriate media off site for disaster recovery purposes. Businesses cannot afford to put their information at risk.

Unfortunately, whether it is an accident-prone user deleting critical documents or spilling liquids on a laptop and destroying their hard drives or a 7.2 magnitude earthquake knocking an entire city offline, information is under constant threat.

Businesses need to employ best practices that promote effective disaster recovery. This starts with being vigilant about backing up files. Organizations need to test file they also test file and volume restores on a frequent basis, and then move backups offsite. Businesses also need to move data offsite in the event of an incident. Businesses today have access to a number of online, inexpensive services that address disaster recovery.



With market research firm IDC forecasting the volume of data to increase by a factor of 4.5 between 2009 and 2012, keeping up with the amount of information will only become more of a challenge. Small businesses may need to solicit a solution

provider to assist in the process to ensure that it is carried out correctly. The investment will more than pay off, providing small business with the assurance that one of their most vital assets is protected, stable, and accessible. ■

-
- i. Nemertes Research, "Benchmark Convergence Reality at Last," November 2004, http://www.nemertes.com/articles/the_virtual_workplace_welcome_to_the_21st_century
 - ii. Computerworld, "Disk Drive Failures 15 Times What Vendors Say," March 2, 2007, http://www.computerworld.com/s/article/9012066/Disk_drive_failures_15_times_what_vendors_say_study_says_
 - iii. Rubicon Consulting Inc "Data Backup for Small and Medium Businesses: Priorities, Current Practices, and Risks," November 24, 2008, http://scm.symantec.com/smb/en/sites/default/files/smb_solution/20002888_ga_wp_databu_smb_priort_currtrprct_risks_p_18403.pdf
 - iv. Smith, David M. and Williams, Michael L., "Data Loss and Hard Drive Failure: Understanding the Causes and Costs," <http://www.deepspar.com/wp-data-loss.html>