

Data Protection & Remote Backup

Preventing Data Loss

For the small to medium enterprise, preventing data loss has become a key component to ensuring business success, for a number of clear reasons. Losing data can cause financial loss, reduce productivity, and increasingly lead to legal entanglements.

Data is perhaps a company's single most important corporate asset. For example most businesses store all customer records, invoices, and banking details digitally within the organisation. If this data were irretrievably lost, a service-orientated business may find difficulty in continuing to trade.

For those able to work from an existing papertrail, the amount of staff time wasted trying to rebuild the systems from a paper trail, or waiting for the systems to be restored before they can get on with their work, can seriously hamper productivity. Today's telephone customers can be less than forgiving when their documentation isn't available instantly, and for some industries, government regulation requires the storage, timely availability to the relevant authorities, and (eventually) destruction of company data. This is no longer an issue for US banks and Sarbanes-Oxley: there are any number of UK and EU requirements that your company may be under – including data protection and industry specific codes of practice.

In order to prevent data loss, a business must have an effective backup practice in place. To be effective, any backup system needs to address three issues of reliability: execution, security, and recovery.

Execution

The most important element to a backup practice is reliability of execution: the backup actually needs to happen, and to be useful, it needs to happen regularly. The frequency of your backups is important: it determines the maximum amount of any data loss - the more frequent, the less data will be lost. Depending on business sector, some companies may be prepared to lose some data: if losing a day's work is acceptable, then a daily backup is adequate, but for some an hour's worth is too much.

Given human frailty, regularly really means 'automatically'. It also needs to be comprehensive both in terms of scope – the type of data, and scale – number of data sources. Documents, applications, and databases are all a necessary part of a company's data infrastructure, and therefore you need backups of all of them. (Consider: with today's constant round of upgrades and patches, installing as opposed to restoring software can be a time-consuming, unproductive business).

Securing data on the company's server is usually insufficient: small businesses' data is spread across a number of user machines. This becomes even more of an issue where a company is located across several geographic sites, or where important data is held on portable devices such as laptops.

Security

Data loss usually comes in two flavours: loss due to physical causes (e.g. fire, flood, theft, as well as the more mundane crashed hard disk), and loss due to human error. Whilst backing up files to a server or to a local tape drive offers some protection against the former, when assessing the risk of physical disaster there is no substitute for off-site

Preventing data loss:

- ¶ Will minimize financial loss,
- ¶ Help with regulatory compliance,
- ¶ Avoid data loss, & thus Reduce productivity loss.

Backup procedures need to:

- ¶ Happen automatically,
- ¶ Happen frequently,
- ¶ Be comprehensive in scope, &
- ¶ Be comprehensive in scale.

Secure backup requires elements:

- ¶ On-site to offset operational risk,
- ¶ Off-site to offset physical risk, &
- ¶ Requires Multiple backup media.

backup. This can be done either manually (usually on an ad hoc basis, or more formerly via a secure courier service), or electronically (via a remote backup system). Backup security is also a question of the reliability of the storage medium itself. Errors can occur when taking backups, and media can deteriorate over time: it's important to ensure that you're not reliant on just a single backup medium.

Recovery

Some companies treat their backup system as 'write-only-memory'. Whilst being glad of the reassurance that having a system in place gives, when things actually do go wrong, there may be no-one available to staff to actually recover lost or damaged files from the backup. Re-keying a document because it's quicker than waiting for IT to recover a file means that any effort put into backup is essentially wasted effort. A backup system needs to ensure that it's easy to recover the correct file, preferably without the need for support staff. If users have confidence in their ability to use a company's backup resource, they are more likely to comply with any requirements that IT staff make of them concerning the backup policy.

Remote Backup

Remote Backup is the process of transmitting your files, via a network connection, to a geographically remote location where they can be stored safely for later retrieval. Remote Backup addresses both the execution and physical security concerns of a company wishing to prevent data loss. Since the process is entirely network-based, the scope for human error (forgetting to take the tapes home on Friday evening) is heavily reduced, and the reliance on third party agents (couriers) is removed.

Why now?

For those companies looking to provide Remote Backup services, the transmission requirements of Remote Backup were a daunting obstacle. In the era of the 28.8k modem, (when a 4Mb file might take 25 minutes to send) remote backup could be done via ISDN or using faster fixed-line connections, but the cost was prohibitive, both for potential service providers, and for many potential customers.

Recently a number of improvements in general infrastructure and technology have made Remote Backup an attractive proposition. Broadband is (practically) ubiquitous, with the UK leading Europe in terms of broadband availability. Most home users now take download speeds of 512Kbit/s for granted, and 2Mb/s are commonplace. Although the majority of broadband connections are via ADSL, thanks to Local Loop Unbundling (LLU), new telecom providers are offering businesses in larger UK cities faster ADSL2+ and SDSL services. So despite ever-increasing file sizes, the transmission needs for Remote Backup are now economically viable.

In terms of secure backup, where the transmission of say financial data is a concern, the hardware associated with secure networking has also fallen in price. ADSL modem

File recovery needs to be:

- ¶ Within end users' capabilities &
- ¶ Possible without IT Support.

"The fact that 40% of businesses suffer a terminal failure as a result of an incident proves that more needs to be done. ... Something like Buncefield, where businesses three miles away were affected, should bring it into focus. ... They need to prepare for the impact on trade, turnover and stock."

Professor Jean-Noel Ezingard, HMC

Buncefield Oil Terminal

The lessons from incidents such as the Buncefield Oil Terminal explosion are two-fold. The first is simple: It can happen to UK companies. A study for the East of England Development Agency (EEDA) study found 25 businesses seriously affected by the blast. Due mainly to the physical damage, roughly 1,500 jobs were relocated, and roughly 80 people lost their jobs due to company closures.

The second is more reassuring. A number of large IT suppliers, despite having their physical buildings destroyed in the blast, were able to return to trading in short order – in the case of Northgate Information Solutions, with just two days of data lost, thanks to a well-implemented manual off-site tape storage.

According to press reports, the fact that it was two days worth of data and not one, was down to poor timing: the previous Saturday's tapes hadn't physically left the site, and were awaiting collection at 7.00am.

routers capable of up to 32 VPN connections are available at high street retail outlets for less than £200; the majority of modern operating systems (Windows, MacOS X, Linux etc.) include direct support for VPN connections.

Despite the disappearance of the cost disadvantages of Remote Backup, for small to medium enterprises and service companies wishing to provide it, the cost advantages remain. By outsourcing the hardware, maintenance, and staffing costs to service companies, small businesses can reap the financial, productivity, and compliance benefits that Remote Backup offers without the capital outlay. Service companies can look to offer Remote Backup to a wider potential market, making it easier to offset these same capital costs. Staffing costs can also be drastically reduced via the use of the latest generation of backup management software.

Using PresSTORE for Remote Backup

PresSTORE is a client-server backup, archive and synchronising package. Several of its key features make it particularly applicable for companies wishing to implement Remote Backup internally, externally, or as a remote backup provider.

Internet Savvy

Firstly, it's designed for the Internet. As well as all data transfer being done over Internet Protocols (IP), the user interface is entirely web-based. Not only can you remotely backup to and from anywhere that the Internet can get to, you can control the process remotely too. For the end-user, this means you can choose to use a service provider who is geographically remote from you, and not subject to say localised flooding. It allows the service provider to address a larger potential audience.

Taking advantage of the lingua franca of standard Internet protocols, PresSTORE further extends a provider's potential client-base by being platform agnostic. The PresSTORE server will run on Unix, Apple Macintosh, and Microsoft Windows. PresSTORE uses a platform-neutral format to store the data, ensuring that platform-specific details (filenames, forks, streams, and attributes) are maintained from backup to restore irrespective of the server and client platforms. For end users this means that they can implement a PresSTORE remote backup now, without switching from their current platform unnecessarily, or indeed, preventing them from switching platforms in the future.

Fast

PresSTORE is fast. Thanks to something called 'Parallel data streaming', a PresSTORE server can read data from a number of different sources, and write the data to several different drives at the same time. Attaching further backup devices to your PresSTORE server, increases the rate at which an end user can backup up to their service provider, or the number

"BT will this week break through the ten million broadband barrier, smashing its initial target of five million connections by the end of 2006. That target, viewed as highly ambitious at the time, was set in April 2002 when there were fewer than 150,000 DSL connections. Broadband availability at that time was 66 per cent whereas it is now more than 99.8 per cent."

BT Press Release, Jan 2nd 2007

"99.6% of UK premises are now connected to a DSL-enabled telephone exchange at 512 kbit/s speeds or higher."

International Communications Market 2006, Offcom

Jargon

ADSL: Asymmetric Digital Subscriber Line: A method for transmitting data over normal telephone cabling (POTS) using high frequencies. The Asymmetric refers to the fact that transmission rates to (upstream) and from (downstream) the service provider are different. In general ADSL offers up to 400 kbps upstream, and 8Mbps downstream. (These are gross rates — this equates to file transfer speeds of roughly 7 Mbps.) Connection speeds are dependent on the physical distance of the user from the telephone exchange.

ADSL2+: By using higher frequencies, ADSL2+ improves data rates to as much as 1Mbps upstream, and 24Mbps downstream.

SDSL: Symmetric Digital Subscriber Line: As the name suggests, this supports the same data rate upstream as downstream. Rates of 2Mbps are possible. For remote backup, this is a preferred option, due to the improved upstream data rate.

VPN: Virtual Private Network: By encrypting the data to be sent and splitting it into IP packets (cryptographic tunnelling), VPN's provide a secure method of connecting a remote user to a network. Thanks to encryption techniques, any data is kept confidential, the two communicating parties can be authenticated, and any data sent is tamper-proofed.

of new clients a service provider can handle without hitting existing users. PresSTORE will even let you utilise additional internet-based servers on demand.

But with backup, speed isn't just about how fast you can save the data. Traditionally, when retrieving files stored on tape, the tape would have to be read at several times during the process in order to locate the correct file – with large storage tapes, this 'seek' time can be an appreciable delay. PresSTORE employs something called "block-direct positioning" – in effect it 'knows' exactly where on a tape to go, eliminating the 'seeking' overhead. Where lightning backup times are critical, PresSTORE's synchronisation tools can be used to synchronise data from the end user's local files server say, to a local hard disk; the backup can then be run from the local hard disk to the remote backup site, without any interference with the local files server.

Reliable

Although PresSTORE is happy to work with hard drives (and optical drives, Network Attached Storage devices, Storage Area Networks etc.), tape drives are its forte. PresSTORE works with a vast number of existing and legacy tape drives and jukeboxes, using a number of techniques for the prevention and correction of tape errors. PresSTORE maintains a record of each individual medium's errors, throughout its life cycle. PresSTORE also offers a block-level cloning feature to produce exact copies of tapes – so if a primary tape fails, the secondary tape is there.

PresSTORE has another trick up its sleeve in the form of an 'indexed referenced backup'. To save time during backup, most backup systems (remote or otherwise) make use of an incremental backup – storing just those files that have changed since the last time a backup was performed. PresSTORE takes this a step further, recording not just the file changes, but the file moves and deletions as well. Not only does this allow a backup to be interrupted at any time, but by combining previous and full backups it's possible to create a full Synthetic Backup, containing a snapshot of your data at any point in time. For end users this means that if disaster strikes, their remote backup supplier can courier them a cloned drive containing their data. For would-be suppliers with existing hardware customers, it's even more compelling: if disaster strikes, we can ship you a replacement server, with your data already in-place.

Simple, Secure & Automatic

For end users, PresSTORE's web interface makes it straightforward to restore files directly to their machines without the need to call on tech. support. For Administrators, PresSTORE's customisable interface means that they control who can do what, and when. For both, PresSTORE handles security; the transfer between sites is secure, since it is the remote PresSTORE server that 'dials out' to the client machines, over a VPN. Scheduling is also handled via the web interface – PresSTORE server can handle the day-to-day operation automatically. And thanks its ability to usefully stop a backup at any point, PresSTORE can even be set to run backups within allocated time-slots.

Conclusion

Remote Backup is something that any data-dependent SME should be investigating. Remote Backup offers a professional solution to key data protection issues: the physical vulnerability of on-site backup, and the reliance on staff (and associated security issues) 'taking home' Friday night backups. As a particular Remote Backup solution, PresSTORE goes further by clearly addressing issues of recovery speed, reliability, legacy systems, platform lock-in, and client usability.

User Scenarios

PresSTORE is an ideal solution for large, multi-user environments such as Magazine or Newspaper production environments characterised by:

- ¶ Multiple sites – PresSTORE's interface makes it easy to manage multiple servers.*
- ¶ Multiple content owners – easy to define user privileges let system administrators keep areas of responsibility separate.*
- ¶ Remote data sources – photojournalists can take advantage of PresSTORE too.*
- ¶ Multiple platforms – written content may be stored on proprietary unix-based systems whilst images and layout are Mac-based. PresStore handles both seamlessly.*
- ¶ Fast changing data – PresSTORE's archive module, makes getting at old stories easier.*

It's also ideal for video production environments, where typically:

- ¶ Huge data requirements – PresSTORE ability to write to multiple back-up devices at once can remove the single-device bottleneck.*
- ¶ Existing SAN & NAS systems – PresSTORE can take advantage of existing hardware*



Enquiries to:
JPY plc, 5 Surbiton Hill Road,
Surbiton,
Surrey, KT6 4TW
Tel: +44 (0) 20 8390 8487

Website: www.jpy.com