



Better backups build better businesses.

21st Century backup technology use expands while old standbys remain viable solutions.



Today's backup needs are orders of magnitude more complex than they were a generation ago. Businesses are generating data at an alarming rate, whether it's social marketing material such as viral media files, purchase orders, proposals, e-commerce database records, incoming e-mail or any of a thousand other digital creations that take place every minute.

As computing power has grown, corporate networks have become more complex as well. Consequently, any business can find its backup strategies fail to keep up with needs, as new technologies and increasing regulatory demands expand the data that must be retained and the speed with which it must be recovered when something goes wrong.

Fortunately, increasingly sophisticated tools can help businesses back up and, more important, use those backup files to recover quickly and efficiently. And new technologies help reduce the size of backups, allowing more affordable storage and quicker access, while reducing overhead.

Backing up in changing times

"Five or 10 years ago, typical medium- or large-sized businesses kept their operational recovery backups in a seven- to 14-day rotation," says Greg Gendron, worldwide tape and disk executive for IBM. "They then did monthly, quarterly and annual backups as their archive.

"That's not tenable anymore," he adds. "Today, they need to change their mindset, as well as their processes and technology, to get archive data out of the backup stream."

"There is more data than ever before to copy, move through the network and store," says Richard Harrison, senior manager of product management in Symantec's information management group. "Just the massive amount of data alone is forcing companies to find ways to get data backed up faster and more securely."

Because businesses are so dependent on technology, data stores and automated processes,

the recovery time objective (RTO) — the amount of time a company can afford for a system or application to be offline — has shrunk dramatically. Likewise, the recovery point objective (RPO) — how much data a business can afford to lose since its last backup — has gotten tighter.

In fact, in some instances, an application's RPO might be zero. For instance, an e-commerce database restored without the transactions completed the moment before it went down means lost sales, upset customers and potentially legal liability.

"In most organizations, expectations for recovery are much greater than they were before because organizations are far more dependent on their IT systems for functioning," says Mathew Lodge, senior director of the Information Management Group at Symantec. "Even e-mail, which used to be considered somewhat trivial, has become mission-critical."

Virtualization has also added to today's backup demands. "Virtualization has created a dramatic change in the number of servers there are and where they're located," continues Lodge. "This has made it more difficult and more complicated to do backup and recovery because now you have a blend of physical and virtual servers."

Today's companies are also challenged with heightened regulatory and legal oversight regarding their information management initiatives. Compliance with regulatory mechanisms, such as Sarbanes-Oxley (SOX) or the Health Insurance Portability and Accountability Act (HIPAA), has drastically increased the data that needs to be saved and the degree of accessibility.

10 backup best practices.

- 1. Use offsite storage wisely.** For the fastest recovery, only the most up-to-date and critical data should be transferred offsite. If you use offsite storage for long-term storage too, make sure to keep them distinct.
- 2. Firewall backup data and archive data.** Streamline your backups by maintaining a clear distinction between archival data and data needed for business continuity.
- 3. Increment, compress, dedupe.** Use every resource at your command to minimize the size of your backup files, both to limit the amount of storage space you need and to speed up disaster recovery.
- 4. Don't forget remote users.** Because they aren't on the network constantly, computers used offsite, like notebooks and teleworkers' PCs, need special attention to make sure their data is captured in your backup system and available to restore.
- 5. Save copies for the least time possible.** Saving old backup data that is available elsewhere wastes storage capacity and slows down recovery.
- 6. Audit your backup system.** Monitor the integrity of not just backups but the storage and network hardware on a regular basis. Do a dry run at regular intervals to ensure your recovery process and your stored data are up to the task.
- 7. Virtual servers require special consideration.** Running backup clients within a virtual machine can affect the performance of the whole virtual environment, and even other VMs on the same server. So watch the resource usage carefully. Dedupe VM disk images to make sure you back up only the data in new backups, not the entire operating system.
- 8. Consolidate.** If you are using systems from several vendors for different aspects of your backup strategy, consider consolidating to a single vendor. Management will be less complex and the chance of error will be lower.
- 9. Start with your pain points.** Build outward from whatever areas are giving you the most trouble. There is no one-size-fits-all solution; every company faces unique challenges.
- 10. Use good security practice.** Apply the same diligence to backup copies of your data that you apply to your original data. This is especially important in the case of offsite storage, which is necessarily out of your immediate control.

Fast forward to today

The complexity of data backup and disaster recovery is largely a matter of scale. For a small to midsize business, backup and continuity is typically only one part of the portfolio of one or a handful of IT professionals.

For smaller organizations or teams with under 50 users, software-based products such as Symantec Backup Exec software or hardware systems such as Seagate's BlackArmor NAS 440 network storage server offer automated system backups on the network at regular intervals, as well as provide the ability to restore whole systems or individual files as needed.

For backing up a couple of servers and a few dozen desktops, a few big hard drives might be sufficient. But that doesn't scale when you need to back up a couple hundred servers and a few thousand desktops, which many enterprises would consider a low estimate of their needs.

Plus, enterprises usually have more at stake when a system or application fails. E-commerce systems, customer relations systems and financial applications might be processing thousands of transactions at any given moment — losing even five minutes' worth of data can be crippling to a large enterprise.

CommVault's Simpana and Symantec's NetBackup provide platforms for backing up across multiple complex enterprise systems, including tools for managing virtual servers and clients, traversing multiple network domains and locations, and reducing the storage space needed without risking data loss.

These systems provide the flexibility needed to deal with enterprise-level backup and recovery tasks. For example, they can take advantage of tiered storage systems that reserve fast but more expensive hard disk or even solid state storage for applications with critical RTO and RPO needs. Less speed-sensitive data is placed on slower drives and long-term archival data is shifted to inexpensive but slow tape.

Likewise, backup frequency can be adjusted on a per-application basis. That way mission-critical systems can be treated to full replication, copying each and every disk write directly to a fast backup disk. Slow-changing data is backed up daily, weekly or even less often, thereby maximizing the use of network and server resources by limiting their use for only the most crucial tasks at any given moment.

Best practices

"When you look at data management and protection, you have to look at it from a broad perspective," says Shawn Aquino, CommVault's senior product marketing manager. "You can either see the forest for the trees, or you can see just the individual trees.

"A lot of organizations have deployed disparate solutions to solve each individual problem," he says. "As a result, they aren't efficient because they have a multistep system for backup, which means you're going to have multiple steps in the restore process as well."

Taking a holistic view of backup can be a challenge, especially in a growing company where needs are changing rapidly. As Aquino notes, most companies tend to look at problems as they arise, which leads to a silo approach to backup — each system or application being treated separately.

Fortunately, many backup and recovery systems now are assembled on a modular architecture. That allows for future expansion as needed while still using existing interfaces. For example, CommVault's software blade model replicates the functions typically added by bringing on new servers, but using the same codebase and interface as already-existing installs.

By planning from a big-picture perspective and making sure a system has enough flexibility to expand to accommodate unexpected situations, enterprises can minimize the number of separate systems and interfaces IT teams need to use. This results in less chance of error when an emergency situation arises in addition to lower costs.

"The biggest expense is staff — upwards of 30 percent of the data protection budget," says Lauren Whitehouse, a senior analyst with Enterprise Strategy Group. "That's why it's important to minimize the intervention of staff in backup and recovery processes."

Minimizing storage

Given the amount of data generated by enterprise applications, another important consideration for both cost and efficiency is how to minimize the amount of storage used. An effective backup strategy needs to carefully prioritize the data being saved so that the most critical data is the easiest to access, while less important data is left to restore in good time, or eliminated from the recovery process altogether.

One way to minimize storage and improve recovery times is to eliminate redundancy in the backup through deduplication. Backing up all the data across a corporate network is inefficient because much of the data is already present in previous backups.

Data deduplication works at the level of bits, examining each block within a file to see if it contains strings of bits that have already been saved. An example might be an image file containing the corporate logo and inserted into every letter, PowerPoint, e-mail signature and internal memo generated by the company.

By eliminating those bits and referring instead to the original copy of the data, deduplication can shave quite a bit of data from initial backups, and even more from subsequent incremental

backups. "Over time, the reduction ratio that we typically see is between 10 and 30 times on average," says Shane Jackson, director of product and channel marketing for EMC DataDomain.

Drawing a clear line between live and archival data is another important consideration. Although various regulatory regimes demand you save copies of a wide variety of corporate data for long periods of time, most of it does not need to be recovered quickly, or accessed often at all.

- **93%:** Businesses that had to use their disaster recovery plans.
- **4 hours:** Average time spent recovering from a disaster.
- **\$287,600:** Average cost of a disaster recovery incident.

SOURCE: Symantec

Having clear archiving policies in place can be essential for business continuity. This is especially important with regard to offsite backup. Because offsite recovery means recovering data over slow Internet connections rather than fast local connections, eliminating as much unnecessary data from offsite backups is crucial.

Deduplication also increases the speed with which offsite backups can be restored. "Data deduplication speeds up offsite backup because you only need to send the unique data," notes Jackson. "Likewise, recovery is faster since the same data can be applied everywhere it exists in the system."

Right tool for the job

Although tape backup seems old-fashioned in a world of 2-terabyte hard drives and

super-fast flash drives, it still has a place in modern backup strategies. Tape remains the cheapest and most durable form of data storage and is eminently well suited to long-term archival storage of noncritical data.

"Our feeling," says Lodge, "is that you should just use the most appropriate technology for your application." Businesses already use information lifecycle management (ILM) to manage their active storage. "The important thing is knowing what a piece of information is and what tier it should be on."

Thus, archival material can go on tape, while more current data can be saved to faster hard drives, such as EMC's Disk Library 3000 and 4000 Series. Similarly, the backup process used should be adapted to the nature of the data being backed up.

Mission-critical applications in constant or near-constant use should be replicated in real time to fast disks. Less-active applications can be backed up with periodic snapshots, such as those produced by CommVault's SnapBackup, which creates frequent recovery points without the overhead of replication. Other enterprise data can be copied in nightly backups, deduped and stored in second-tier storage until policies allow it to be either deleted or archived.

The benefits of a holistic view to backup are clear. Reduced management time, lower storage needs and a simplified interface mean less staff time devoted to running the system, and that means lower operational costs. ♦

Contact your **CDW account manager** for information about how to consolidate storage and free up IT resources.