

Choosing the right backup solution

There are a range of products and technologies available to meet your RTO and RPO requirements.

By Adam Fore



ALL DATA-PROTECTION STRATEGIES are subject to the same challenges and outside pressures. The strategies that you choose must take into account these considerations. To start, no thriving organization is exempt from the pressures of managing and storing ever-increasing amounts of data. Your data-protection strategy needs to be scalable to accommodate explosive data growth.

All global operations that must be up and running 24x7 are confronting shrinking backup windows. It is not uncommon for full backups to take an entire weekend and for an incremental backup to take all night. At the same time, backups are taking longer to complete and backup windows are getting smaller. Rising costs for backup management are the direct result of the data explosion.

In a recent survey, International Data Corp. reports that, on average, 30% of the IT storage budget goes to data protection. Getting data off primary storage and onto backup devices without creating performance bottlenecks can be a big challenge when you have tens to hundreds of terabytes to back up.

Compliance requirements and legal discovery demands are growing. Data must be secure, unalterable, and immediately accessible in order to comply with regulations such as the Federal Rules of Civil Procedure in the U.S., the European Union Data Protection Directive, and Japan's Financial Instruments and Exchange Law (J SOX).

The recovery point continuum

A typical backup architecture needs to provide for a wide variety of environments, including NAS, SAN, and direct-attached storage; data center and remote offices; a variety of operating systems; and a range of recovery point objectives (RPOs) and recovery time objectives

(RTOs). When choosing a backup strategy, one must view it from an RTO framework. You should consider a range of solutions that consider your RPO and RTO for your backup data (see figure).

Choosing a backup solution

In choosing a data-protection solution, you must understand the varying protection needs for the different applications and data in your environment.

For each classification of data, consider these questions for both operational recovery and disaster recovery:

- **How much data loss is acceptable in the event of a failure?**

This helps define the protection frequency. The maximum time requirement between a failure event and a point-in-time to recover from is called the RPO.

- **How fast do you need to recover the data after a failure?**

The required time it takes to make data accessible after a failure is called the RTO.

- **What is the operational backup window? How much time do you have to perform backups?**

This defines the backup performance or backup technology required.

- **What is the budget you need to work within?**

- **How open are you to changing your existing backup infrastructure or processes to address the backup issues?**

- **Do you plan to continue using tape for on-site and off-site retention?**

With advances in data protection it may be possible to consolidate tape usage at a single site or extend the current tape investment.

Your answers to these questions will help you identify suitable data-protection technologies.

Backup to tape


Backup to tape is currently the most widely deployed solution for data protection and is at the lower end of the RPO/RTO continuum. Tape-based copies of data may still serve multiple purposes, including operational recovery, disaster recovery, and archive storage.

Tape backup is characterized by low media cost, portable media, and the fact that it is offline, un-powered storage. Backup to tape has a long history and is so tightly integrated with backup applications that much of the backup software in use today has been engineered specifically for tape backup. However, relying solely on tape for diverse needs has become an increasingly risky proposition.

Due to the amount of time to restore data from tape,

operational recovery and disaster recovery pose significant challenges for tape backup. Data on tape can only be accessed sequentially, which significantly impacts recovery times. Another major problem is performance: Bottlenecks getting data off primary storage and onto the backup device are inevitable when you have tens to hundreds of terabytes to back up.

Tape requires frequent, full point-in-time copies of data for effective recovery. Tape's role may change even though a recent IDC survey reported that 25% of



respondents indicated that tape will be replaced. Tape may eventually be used only for archival and off-site storage purposes.

Backup to disk

The first generation of disk-based backup is known as "backup to conventional disk." This type of disk backup, in which backups are streamed to a conventional disk subsystem usually filled with low-cost SATA drives, has been around for a long time.

Traditionally, high acquisition and management costs have limited the use of conventional disk to only the most recent backup copies. The advantages of conventional disk are improved recovery times and protection against media failures by means of RAID technology. Higher aggregated throughput can also be achieved using the ability of conventional disk to accept multiple backup streams without the need for multiplexing data on tape media.

Data on conventional disk can be accessed randomly, significantly increasing the speed of single file, directory, or volume recovery. Yet, acquisition costs for conventional disk remain high due to the lack of

advanced storage-efficient technologies (e.g., data de-duplication).

Backup to virtual tape libraries

As we move along the RTO continuum to increasingly stringent requirements, we come to more-advanced methods of disk-to-disk (D2D) backup. First of these is virtual tape libraries (VTLs) and then VTLs that offer advanced D2D backup features. These are easier to integrate because they already have existing backup applications and are non-disruptive to both backup administrators and existing processes. Some VTLs also facilitate de-staging data from the VTL to physical tape to accommodate off-site tape vaulting for disaster-recovery purposes or long-term archiving. Of course, another method with this option is to clone off-site. Enterprises for which VTLs are the best option exhibit one or more of the following characteristics:

- Preserve their investments in existing backup software and tape systems;
- Use disk and tape in a tiered backup strategy; and
- Improve an existing tape backup environment that is difficult to change.

To meet the backup window, it is essential that the VTL be capable of delivering the aggregate write performance required both today and over the planned life of the system (typically three to five years). To be cost-effective, the VTL must also provide capacity-efficient storage without sacrificing performance, using either standard compression, data de-duplication, or a combination of the two. To be easily manageable, the VTL should be certified for use with leading backup applications and designed to be managed by a backup administrator.

It is important to consider the impact to your existing tape infrastructure. Ideally, you should search for a system that optimizes tape creation performance without increasing the consumption of physical tape media. Most important, you should find a VTL solution that offers seamless recovery so that off-site cartridges are easily read back.

Snapshots with replication

As we move further up the RTO continuum, we come to disk-optimized solutions where there are many data-protection services that capture changes to data and

write them to a separate storage location. There are multiple methods for capturing these changes, involving different technologies that serve different needs; they may capture data at the file level, block level, or application level. This capture may be done on a point-in-time manner or continuously.

SNIA definition of snapshot: A fully usable copy of a defined collection of data that contains an image of the data as it appeared at the point in time at which the copy was initiated. A snapshot may be either a *duplicate* or a *replicate* of the data it represents.

Snapshot technology is a cost-effective way of providing a solution to meet a small retention time objective. Remote replication adds a disaster-recovery solution without impacting the current application environment.

Snapshots make frequent point-in-time copies to provide granular restore points and fast recovery times. By providing recovery points as frequently as every minute and by being very network- and storage-efficient, snapshots are very cost-effective for a wide range of applications and data.

A snapshot and replication solution can shorten backup times from many hours to a matter of minutes without impacting the backup process or application performance. Like many of the newer D2D technologies, snapshots provide more-efficient storage through inherent de-duplication of data—only changed blocks of data are transmitted and stored to provide what is essentially a full backup.

Backup data can be served directly to the backup system. Data can be restored from moments before a failure or disaster. Restores can happen quickly because end users can restore their own data. Snapshot backups can be used for other business needs, such as test and development or decision support.

It is important that a snapshot with replication solution is integrated with your application to ensure the resulting snapshot backups are recoverable by your application. Look for a solution that meets your requirements first and then look for one that provides efficient use

of storage capacity. For many environments, an ideal solution will be manageable by your existing backup tools, scalable to hundreds of terabytes, integrated with leading enterprise applications, and work in both data center and remote office environments.

Some solutions can facilitate faster recovery by storing data in native formats so that it can be recovered directly by authorized end users and administrators.

Continuous data protection (CDP)

SNIA definition of CDP: Continuous data protection (CDP) is a data-protection service that captures changes to data to a separate storage location. There are multiple methods for capturing the continuous changes involving different technologies that serve different needs. CDP-based solutions can provide fine granularities of restorable objects ranging from crash-consistent images to logical objects such as files, mailboxes, messages, etc.

Another data-protection service that captures changes to data on a separate storage device is CDP. CDP records every

Backup strategies do not take into account the security and privacy of the data itself.

write with no disruptions to the systems being protected. There is no backup window and no scheduling to manage. Recovery to any point in time may be appropriate for high-end applications with a near-zero RPO, and for environments where IT involvement in the data-protection process needs to be minimized.

RTO can also be improved with CDP. Recoveries take advantage of high-performance disk, may be at very fine levels

of granularity, and may provide advanced search functionality. Some, although not all, CDP solutions lack application coherence and consistency. This is not a challenge that is unique to CDP; it can also be a problem with snapshots. CDP systems that do not offer application consistency may roll back to any point in time, but they cannot guarantee that the application was in a consistent state at the chosen recovery time. It is important, when implementing a CDP solution, to assess the disk and network capacity requirements. Because CDP captures every change to data, you will also need to manage your retention policies in line with available disk capacity.

Space efficiencies

SNIA definition of compression: The process of encoding data to reduce its size. Lossy compression (i.e., compression using a technique in which a portion of the original information is lost) is acceptable for some forms of data (e.g., digital images) in some applications, but for most IT applications, lossless compression (i.e., compression using a technique that preserves the entire content of the original data, and from which the original data can be reconstructed exactly) is required.

It is important to discuss in more depth the key characteristics that contribute to space or capacity efficiencies.

Compression

Data compression encodes information to use fewer bits. This reduces the consumption of resources such as disk storage and network bandwidth to accomplish greater efficiencies. Generally, software-based compression uses server CPU cycles to achieve its typical 2:1 compression ratio. Hardware-based compression mitigates this performance issue by offloading the compression function to a dedicated chip. The actual compression ratio depends on the type of data being compressed. JPEG files typically don't compress at all, while Microsoft Office documents can be compressed to a tenth of their size. Compression significantly increases the efficiency of some D2D backup solutions.

De-duplication

De-duplication is a term used to identify data-reduction technologies that eliminate the redundant copying of data at some level of granularity. Non-optimized backup methods generate duplicate data as part of the backup process—every full backup is a duplication of a data set. De-duplication technology eliminates this redundancy without affecting data recovery. By discarding duplicate data strings and updating data pointers, de-duplication reduces disk capacity requirements and enables you to store more backups more cost-effectively. All data can be safely recovered in its original format from the disk volume. De-duplication technology makes D2D backup solutions more affordable for a broader range of data. In many cases, de-duplication can work in conjunction with compression to yield even greater resource savings.

There are three qualities that you should expect in a de-duplication technology:

- More duplicate data will be found when the data blocks being de-duplicated are larger;
- The process of de-duplication should not affect backup or recovery performance; and
- Enterprise-class reliability and media protection features are a requirement, because a single data block error or media failure can affect multiple backups.

If your data is encrypted, compression and de-duplication are of marginal value.

Extending backup strategies and value

D2D backup strategies continue to evolve by integrating more-specialized technologies, thus bringing more value to your business. For example, making your backup data more discoverable makes it easier to meet legal requirements. Indexing and accessing data can be much faster

when backups are on disk. Disk backup makes it possible to quickly search for and locate individual files on storage systems throughout the network—an ability that is invaluable when carrying out an e-discovery search in response to litigation. Classifying data ensures a better data-protection practice. You can set up different data-protection efforts depending on the type of data.

Today's backup technologies go a long way toward protecting your data. However, they do not take into account the security and privacy of the data itself. By nature, backup procedures introduce additional threats to stored data: With each additional distributed copy of cleartext data, you increase the risk of unauthorized access. By encrypting data before it is ever written to disk or tape, you can ensure only authorized people are able to read data. Your data can be fully protected against unauthorized access if a disk or tape is lost or stolen.

Conclusion

Choosing the right backup solution starts with a careful examination of your data classifications and availability needs. Along with this, many D2D backup strategies are coordinated with disaster-recovery strategies. The scope of data protection continues to mature with demands for insurance of data loss while leveraging additional value of backup data. □

Adam Fore is a member of the SNIA Data Management Forum (DMF), and a senior manager, enterprise data-management products, with Network Appliance.

BULK REPRINTS

Do you need bulk reprints of an *InfoStor* article? If so, contact **Andy Speter**
 Tel: (212) 221-9595 x250
 Fax: (212) 221-9195
andy@parsintl.com

The benefits of data footprint reduction continued from p. 33

may occur as a result of using SIS-based data footprint reduction techniques during large-scale bulk data restorations.

You may want to consider a data footprint reduction strategy that combines various technologies to address specific applications as well as your overall environment, including online, nearline backup, and offline archiving.

reduction technologies for future use, including archiving with data discovery (indexing, e-discovery), consider leveraging appliance-based compression technology to maximize the capacity of existing storage resources for online, backup, and archiving, in conjunction with other data footprint reduction capabilities;

of storage, business applications, and other functions;

- Data archiving should be an ongoing process that is integrated into your business and IT resource management functions, as opposed to being an intermittent event to free up IT resources; and
- Get a handle on your data footprint and its impact on your environment using analysis tools and/or assessment services. Develop a holistic approach to managing your growing data footprint. Look beyond storage hardware costs, and factor in software license and maintenance costs, as well as power, cooling, and staff management time.

There are several different techniques that can be used individually to address specific data footprint reduction issues, or those techniques can be used in various combinations to implement a more comprehensive and effective data footprint reduction strategy. The benefit of a

broader, more-holistic, data footprint reduction strategy is to address your overall environment, including all applications that generate and use data as well as overhead functions that impact your data footprint.

Reducing your data footprint has many benefits, including maximizing the usage of your IT infrastructure resources such as power and cooling, storage capacity, and network bandwidth, while enhancing application service delivery in the form of timely backup, BC/DR, performance, and availability.

Look to combine technologies and techniques to address your various data footprint challenges, and to maximize your IT resources while reducing management costs and complexity. □

Greg Schulz is founder and senior analyst of the StorageIO Group (www.storageio.com). He is also author of the book, *Resilient Storage Networks—Designing Flexible Scalable Data Infrastructures* (Elsevier Digital Press, 2004). He can be contacted at greg@storageio.com.

You may want to consider a data footprint strategy that combines various technologies to address specific applications.

Following are some general recommendations and suggestions to help address your growing data footprint, all of which depend on the size and scope of your particular environment, applications, and service requirements.

- If you are evaluating data footprint

- Maximize use of your existing IT resources without introducing complexity and costs associated with added management and interoperability headaches. Look for solutions that complement your environment and are transparent across different tiers