

DEMYSTIFYING DATA PROTECTION TECHNOLOGY

Running a business is frightening work. In addition to normal operational challenges, you must comply with various regulations or face legal penalties. News stories about the 80% businesses that never reopen after closing due to a loss to criminal actions, natural disasters, unprovoked virus attacks, or human error are everywhere. On top of all this, customers now demand 24 hour access to your business and trading partners increasingly are asking for performance guarantees.

Data Protection Concepts

When you ask your IT staff (assuming you even have an IT staff) about what steps you should take to protect your business, they respond with “techno-speak” about disk-to-tape backup, disaster recovery, data protection, and business continuity. What do these terms really mean and what are your responsibilities, legal liabilities and business options?

This paper examines each of these concepts, discusses how and when to apply these technologies and explain the role that they play in crafting a comprehensive data protection strategy.

The term “data protection” is a catchall phrase that applies to a wide variety of technologies and operational procedures.

Data protection consists of three distinct but related concepts:

- Data Redundancy
- Data Archiving, and
- Data Replication

Each of these will be discussed, but first lets build a common vocabulary.

Three Important Concepts

Two terms at the center of the concept of data protection are *Recovery Point Objective* (RPO) and *Recovery Time Objective* (RTO)¹.

When you establish a Recovery Point Objective, you set a limit on the amount of information your organization can lose and still operate. Beyond this point, you loss the ability to recover various business transactions and you must declare a “disaster.”

A Recovery Time Objective is a measure of how long a process can be out of operation before it has a fatal impact on the business.

RPOs and RTOs are generally expressed in units of time since this is the easiest way to measure the impact of a severe business interruption on your operations. For example the accounting receivable function might have an RPO of one week and an RTO of ten days – depending when in the billing cycle the disaster occurred.

¹ Precise definitions of these terms can be found in the “Glossary or Terms” section of the Disaster Recovery Institute International website (<http://www.drii.org/displaycommon.cfm?an=3>).



It is important to note that Recovery Time Objectives assume that the information utilized by the process is available and meet certain “freshness” goals. For this reason the RPO for a specific process generally (but not always) have the same or shorter timeline than their RTO counterpart.

Together these terms provide the basis for deciding what level of data and process protection an organization needs in order to stay in business.

With this information, management can weight the level of risk they will tolerate and determine the level of investment they want to make to achieve their recovery objectives. Given a budget, most business continuity planners can develop a list of strategies they should deploy in order to meet these goals.

Different Requirements for Different Functions

RPOs and RTOs differ dramatically by application and department. Whereas stock traders can not sustain the loss of even a few seconds of data, the same firm’s purchasing department might be able to operate for several days without access to certain records such as previous contracts. For the accounting department, as long as the information is available by the next billing cycle the department can continue its function even if access to the accounts receivable list is unavailable for a week or more.

This same variability holds true for each business process. Some functions might have an RTO of zero, as in the case of an airplane flight control system; to several weeks for access to employment applications after a fire has decimated a business.

When calculating RPOs and RTOs it is customary to also estimate the financial implication that these outages would have on the business. While a crude form of financial justification, these estimates do help form the basis for gauging the amount of investment to make in various protection schemes.

Recently a new measurement concept has been proposed that attempts to track the exposure an organization has based on regulatory and compliance reporting rules. Dubbed RCO for Recovery to Compliance Objective, this is also time-based metric and measures how long an organization has before it must report its reduced capabilities to a regulatory agency.

Online, Offline and Nearline

The most common data replication technique involves copying files from one or more disks to magnetic tape (s) for storage and then filing these tapes away either at the same site or at a remote location.

The term “online” storage is used for data that is directly accessible by a computer program without the need for intervention by either a machine or an individual.

Information that is copied onto some media that is then removed from the computer system is referred to as being “offline.”



A third concept is “nearline” storage. If the actual tape media remains resident in a device that can automatically reload it when the data is called for by data an application, the information is referred to as being in a “nearline” state. Many high-end tape and optical storage devices (CDs and DVDs) use loaders such as jukeboxes or carousels to automate the retrieval of the offline media – thus creating nearline environments. Some systems go so far as to bar code the media to facilitate the tracking the physical media location for better control and security.

“Media management” is the term for this process of tracking the physical location of removable media.

Tape versus Disk

Disk to tape (D2T) backup was popularized several decades ago when the price of magnetic disk storage was very high compared to tape. Over time the cost of disk drive units has dropped significantly and this cost imbalance no longer exists. Today, it is a fallacy to believe that tape backup is an economical alternative to disk backup especially given the fragility of tape. Optical storage continues to have a legitimate roll in data protection because of the portability, durability and tamper-proof nature of the media, but even these characteristics are being replaced by magnetic storage technologies such as “thumb drives” and portable disks.

The other terms you will encounter when discussing data protection are D2D, which stands for disk to disk backup and D2D2O, which is a relatively recent term. The “O” refers to Offsite which means a different physical location. When using a D2D2O strategy it is very possible for information to be *offsite* while still *online*.

Data Redundancy

Making a copy of a vital record or critical information is the oldest and simplest form of data protection. The computer term for this process is “data redundancy.” This protection strategy is useful in situations where there is concern about accidental or purposeful destruction of the original data. While a duplicate of the original is prepared, this term does not imply that the information is safe, secure or stored in remote physical location.

Network backup programs, which run according to a fixed schedule, are excellent examples of this rudimentary form of data protection. Once you understand the RPO for particular information set, you can make a decision about how frequently to run a backup session. For example, if the RPO for the information used by your customer service desk is four hours, then every four hours or less, you should run some form of data duplication backup.

There are some significant shortcomings to the D2T approach. Although D2T provides a low cost data protection mechanism there are several disadvantages that make this approach unattractive. For example, magnetic tape media is very susceptible to damage from heat, electro-magnetic interference and physical handling. Second, it is not easy to test the quality of the information written to tape to insure you have an accurate recording.



This lack of quality control casts a shadow over the confidence one has in tape-based backups. This issue has been the topic of study of many well-known research projects, some of which report that up to 40% of tape backups have serious errors that could result in the loss of data.

Another key shortcoming is that information stored offline on tape is effectively invisible and therefore much less manageable and poses additional security concerns.

The D2T model has historically been viewed as a low cost way to obtain data protection. Declining online disk storage costs are doing much to change this perception. Add to this the cost of not being able to recover even once because of lost or misplaced volumes, damaged media, or incomplete file and the economic incentive to use removable media disappears.

Despite its disadvantages, tape-based backup remains the most common form of data protection used in data centers and even local area networks today – although this is changing.

Fueled by awareness of the shortcomings of magnetic tape, many organizations are substituting optical disk media (CD-ROMs and DVD disks) for removable magnetic media in an effort to find a more durable media and restore confidence in their backup strategy. Unfortunately, while optical media does offer some physical advantages over magnetic tape, the information is still stored offline and movement of the media volumes requires human intervention, which drives up the cost.

A New Approach

There is an industry-wide trend to replace the disk to tape model with a new form of backup called the disk to disk model. The D2D approach offers quick restore times², reduced handling issues, and improved reliability. Many organizations cost justify the movement to a D2D model based on several factors including: the increased control and data visibility online storage offers, the elimination of labor costs associated with media handling, cost saving from the elimination of offsite shipping and handling, and the peace of mind that comes from knowing that the copied information is actually stored in a correct and useable format.

Disk array technology is one category of the D2D data protection that has become increasingly popular. Also known as RAID, this technique is generally applied on a local level, often within the computer unit itself. The reduced size and power requirements of today's disk drives make it possible for even laptops to support RAID Level 1 configurations.

An acronym for *redundant arrays of inexpensive disks*, RAID-1 technology creates a mirror image of one disk on a second one located on the same system. In the event of a

² The key benefit of making a copy of data files is the ability to recover the information they contain in the event of a loss of the original copy. This recovery process is known as a "restore" and it is closely related to the concepts of RPO and RTO.



failure of one disk, the second one can carry on the assigned task and once the failed drive is replaced, automatically rebuild a fully copy of itself on the new media. RAID technology is one way of offering a RPO of zero, and is appropriate for certain types of information.

Unfortunately, even the wide-scale use of RAID technology is not sufficient to protect a business against data loss since the “data copy” is still stored at the same physical location as the primary systems. A fire or other destructive event can permanently stop the retrieval of the critical (and not so critical) data.

Virtual Tape

Despite declining cost and the obvious benefits of online backup, there are still costs involved in implementing this approach – especially in redesigning and recoding many software applications that were designed to work with magnetic tape, not disks.

This explains why many backup programs are now promoting the concept of *virtual tape emulation*. In effect, these backup programs allow a D2D operation to appear to be a T2D one to the software applications using them. This emulation strategy allows organizations to avoid the danger and expense of modifying many of their applications which were designed to work with tape while providing the speed and reliability of magnetic disks.

This strategy represents a “brute force” approach to data protection since the majority of application packages copy virtually everything from the host system to the “target” and *do not* differentiate between vital records and casual files. However, new products from companies like Asigra are changing the face of the data protection market.

Continuous Data Protection

Asigra and a few other organizations have introduced the concept of *Continuous Data Protection* (CDP) which tracks changes at the block level and ensures that a near-real time backup of changing files takes place. Combined with intelligent algorithms that eliminate data duplication and securely compress data to a fraction of its original size, these systems address the challenge of moving large blocks of data to-and-from a remote electronic vault.

Today it is feasible for even a large data set to be safely and securely replicated in a location thousands of miles away from the business site. The high speed restoration of a small number of files is possible through advanced networking techniques. In the event of a large scale loss of data, other techniques including access to a *proxy storage server*³ or the use of *mobile electronic vaults* can provide an acceptable level of data access while the primary site is restored.

³ In this context, the term *proxy storage server* means a remote data store that is accessible in a secure manner over a public or private network and can provide access to the lost data while in parallel a local data store is rebuilt.



Solving A Different Problem

Data Archiving represents a different approach to data protection that is usually reserved for important but unchanging files. Examples of this class of static files are employment records of terminated individuals, past financial records and various government filings such as patent applications.

Unlike backup, archiving does imply that a value judgment about the importance of the information being stored. In many cases, the archived file is the only remaining copy of the data since for many organizations, archiving implies removal of the primary version of the data.

As a practical matter, it is prudent to make at least one copy of all archived information. Applying this approach, we see a combined data protection strategy that identifies and isolates static but important files, and enhances the “retrievability” of a record or file by creating a saved copy. Having at least one copy of the archived file is a risk reducing measure that guards against permanent data loss due to damage or destruction of the primary file.

The popularity of data archiving is increasing in response to government and industry record keeping regulations. This is also an area where magnetic tape and optical media are frequently used as the storage medium.

Advanced Data Protection

Data Replication is the most advanced form of data protection because it combines a number of the advantages of continuous backup and archiving, while utilizing geographical diversity as an additional element of protection. Once viewed as a technically demanding application recent advances make data replication a viable choice for even small institutions.

When properly implemented, data replication procedures move copies of important documents to one or more offsite, but online storage locations. The files at these locations are accessible and remain synchronized with the changing data automatically.

For processes with a long RTO⁴ periods (i.e.: measured in days or weeks) there is a temptation to use magnetic tape as a storage media and keep the information offline as is done with paper or other hardcopy records. Before committing to such an action it is important to remember that the same shortcomings that made use of offline, removable media unattractive for backups apply with equal weight to data replication.

Data replication can operate in two modes: asynchronously or according to a defined schedule. Asynchronous data replication implies the updating of files on an ad hoc/ as needed/ as occurs basis. Scheduled updates occur at defined times, generally during off hours when there is little or no competition for the firm’s network connection.

⁴ Some regulations allow five or more days for the retrieval of certain records. In legal discovery cases the timeframe could be longer.



Both modes of operation require the creation of an initial copy or “seeding” of the files being protected (also known as a “full” backup). At this point, data replication has much in common with backup, but this is where the similarity ends. Unlike backups that sequentially copy the daily changes (incremental backups) to every file and periodically (weekly?) retake full backups of all the files; data replication keeps all files current and up-to-date, just as if they were the primary version of the data. There is no need to perform additional full or incremental backups since the primary and secondary file(s) are always synchronized and current. This approach is also very bandwidth efficient.

Sometimes immediate synchronize of the primary and secondary files is not desirable. Human error or other factors might incorrectly cause the deletion of valuable information. Advanced systems can create a policy that automatically duplicates any replicated file prior to applying certain types of updates such as data deletion commands. In such situations, one version of the file is updated and synchronized, while the other copy is left untouched. This procedure can be repeated innumerable times and the unaltered versions retained as long as desired.

The main benefit of data replication comes into play when the copied data is located at different site(s) geographically separated by a significant amount. This concept of geographical diversity reduces the risk inherent in having all instances of your information located in a single spot. If a fire or some other wide-scale disaster occurs, copies of your selected records are still safe and accessible because they are housed in a remote location. Properly implemented, an asynchronously run data replication model can be used to support a RPO of near zero.

The obvious question is, “Why not store everything online and remote?” The answer to this question has more to do with technical limitation of available data transmission lines, than any business reason. Given a large volume of data, the amount of time needed to move information across a standard DSL or T1 line can become prohibitive.

Process Protection Strategies

This paper would not be complete without mentioning the *process protection* technologies that complement the *data protection* technologies discussed above.

If the RTO for a process is long, then the appropriate recovery strategy may be to order a replacement system and rebuild the environment. There are a number of “quick ship” programs available from popular vendors that will ensure that a suitably configured computer is delivered to the location of your choice in a specified timeframe. Another option is to “reserve” a set number of computer systems that are already installed at an alternative site. Organizations such as Sungard and IBM offer such products.

In both these instances, some rebuilding and customization is to be expected.

For processes with short RTOs, then a variety of products exist that allow the configuration of computer environments that can be put into operations after a brief



period having reconfigured themselves. The term-of-art for this class of systems is High Availability.

Some computers are configured with specialized software that allows them to “ride through” the failure of one of more components with little or no impact on operations. These systems are referred to as being *Fault Tolerant* and differ from high availability systems in that they operate continuously through the use of redundant components without even a few milliseconds of interruption.

Just as data replication adds an element of geographical diversity to the issue of data protection, “application mirroring” brings this same advantage to the area of process protection. Mirrored sites are one where critical systems and subsystems are duplicated and can be put into service quickly. If the equipment sits ideal until activated in the event of a business interruption, the site is referred to as a “cold standby” (i.e.: the equipment is not turned on). If the equipment is installed and running but not operating under a workload, it is referred to as “warm standby” location. If the site is fully function and actually performing useful work, it is referred to as a “hot standby” location.

Putting It All Together

In this white paper, we have examined and discussed many different data storage concepts and technologies. To summarize, when developing a data protection strategy it is essential that you identify your critical business processes, the data that these processes use or rely on in order to operate and any other vital records that your firm may need access to in the event of an emergency.

Next, you develop a Recovery Time Objective (RTO) for each of these processes and the data associated with each process. You must also establish and document a Recovery Point Objective for the data so that you understand how current your information must be if you are to return to operations. Having established these two figures you are now in a position to determine what type of data protection strategy is justified.

Backup is an excellent data protection strategy when the goal is to create a copy of large amounts of non-prioritized data. Given the advantages of keeping backup data online and easily accessible, the D2D model is the preferred way of implementing this data protection strategy.

Data archiving is ideally suited for static files that must be retained for an extended time-period. Once again, keeping these files online and visible insures that they are correct and accessible when needed.

Data replication is an advanced technology that combines many of the advantages of other data protection strategies with geographically remote online storage and very efficient bandwidth utilization.

Data protection technologies can seem intimidating at first. However, once you demystify the basic terminology, even a novice can design data and process protection plans that will keep your business running no matter how sever the operational issue.



About the Author

Don Byrne has spent over 30 years in the High Tech industry with concentrations in storage technology, business continuity, and enterprise software. He has worked as a venture capitalist, a corporate development executive handling mergers and acquisitions, a “*turn around*” CEO, and is the cofounder of five companies. His latest, North River Solutions, offers consulting and business development services. HE holds degrees in Mathematics, Philosophy and International Marketing. He is completing his PhD in Organizational Behavior and is the author of a forthcoming book on business continuity due out in winter of 2007.