

# Data Protection Scenarios and Solutions

An Educational White Paper

By Dave Therrien  
CTO of ExaGrid Systems



## Introduction

There are many different solutions that must be employed today to provide proper data protection for business data. Data loss can occur when users or applications accidentally delete one or more files. Data needs to be protected from virus infection. Data must be restored or recreated when a single disk drive or an entire disk subsystem fails. Finally, disk data needs to be recovered when one or more systems are damaged or lost after a site disaster. To prevent data loss, there are a number of independent data protection solutions that are being employed by IT organizations today, including:

- **Disk mirroring**
- **Parity-protected disk**
- **Weekend and nightly backups**
- **Disk-based weekend and nightly backups**
- **Intraday backups**
- **Remote office backups**
- **Offsite data protection**
- **Disaster recovery**

This paper will review these data protection solutions and discuss the role each plays in formulating a complete IT data protection strategy.

## Disk Mirroring

Disk mirroring, sometimes referred to as RAID 1, produces a bit-for-bit copy of data from a primary disk drive to a secondary disk drive. If either disk drive fails, the other disk drive provides continued access to all data. RAID 1 is typically performed within a single localized disk subsystem.

Disk mirroring is often employed with the primary and secondary disk drive separated by a campus or metropolitan distance. Data can be replicated synchronously or asynchronously.

- **Synchronous** mirroring keeps both the local and remote disk up to date on every write request, which can cause reduced overall application performance. To minimize this impact of disk mirroring on application performance, high bandwidth, low-latency network connections like Fibre Channel, DWDM, or CWDM are employed and are limited to campus/metro distances.
- **Asynchronous** mirroring allows multiple writes to be queued for replication to the remote site. In the event of a failure of the local disk drive, some number of write requests may not have been completed to the remote drive, so the application may lose some of the recent history when it fails over to the remote drive.

Disk mirroring is an ideal solution for applications that demand transactional high availability and transparent failover either when the local disk subsystem fails or to expedite recovery after a site disaster.

Disk mirroring does not protect against accidentally deleted or corrupted files because whatever is written to the primary disk is immediately copied to its mirrored disk. Deleted files on the primary disk are instantly deleted on the mirrored disk, and files corrupted on the primary disk are instantly corrupted on the mirrored disk. Therefore, mirroring does not eliminate the need for periodic backups of data to another storage medium like magnetic tape. Periodic backups of data create historical versions so that, when a file is deleted or corrupted on the mirrored drive pair, a copy can be restored from a point in time before the deletion or corruption occurred.

## Parity-Protected Disk

Parity-protected disk subsystems are a more cost-effective way to protect from disk drive failures than disk mirroring.

The most popular form of parity-protected disk storage is RAID 5, which effectively adds an extra "parity" drive to a set of "data drives" to allow the total set of drives to continue to deliver data even when a single drive fails. For example, a 4+1 RAID 5 set would have four data drives and a parity drive. In this example, only 25 percent more disk storage capacity is required to protect the data drives. In contrast, disk mirroring requires 100 percent more disk storage capacity to protect data.

Some RAID 5 disk sets may also include another “hot-spare” disk drive that is used to immediately start the rebuild process when any of the other RAID 5 drives fail. Alternatively, instead of a hot spare, when a disk drive fails in the set, it can be removed and replaced with a new disk drive, but this can lengthen the time to bring the RAID 5 disk set back to full protection status.

If two drives fail simultaneously in a RAID 5 disk set or a second drive fails while the first drive is being rebuilt using the spare, all data is lost and must be restored, typically from weekend and nightly backup tapes. With disk drive capacities growing even beyond the 500GB range, the time it takes to completely rebuild a failed drive can take many hours, and even days if performed as a background operation. This extended rebuild period creates a greater probability that a second drive could fail causing all data in the RAID 5 array to be lost.

RAID 6 is becoming increasingly popular as an improved parity-protection option for magnetic disk subsystems. For slightly more physical disk storage capacity than RAID 5, a RAID 6 array can continue to operate even when up to two drives in the array have failed. With RAID 6, it takes the equivalent of two drives per array to maintain parity data.

Like mirrored disk systems, RAID 5 and RAID 6 disk subsystems do not guard against deleted or corrupted files and therefore require nightly backups in order to keep a day-to-day history for recovery from deleted or corrupted files.

There are many other configurations of parity protected RAID disk storage. For instance, RAID 51 provides two RAID 5 arrays that are synchronously or asynchronously mirrored either locally or at two different sites.

## Weekend and Nightly Backups

There are three basic types of weekend and nightly backups:

1. **Full backups** – all primary storage data is written to a set of tapes. If a full backup is performed on 10TB of primary data, all 10TB will be written to a set of ten tapes every weekend. Full backups are typically run for all data on all servers during the weekend. In addition, certain applications like databases and email data stores might also be backed up nightly as a full backup. Since the full backup process consumes a tremendous amount of network bandwidth and client system processing and storage cycles, full backups are typically performed on the weekend to minimize the performance impact on users and applications.
2. **Differential backups** – all primary storage data that has changed since the last full backup are written to a set of tapes. If a full backup was completed on a Saturday night, a Wednesday night differential backup would protect all of the files that were created and changed on Sunday, Monday, Tuesday and Wednesday. If differential backups were run every night, the size of backups and the time to perform backups would become too great, so typically, differentials are run once or twice in the middle of the week. The advantage of running a differential backup is reduced time to restore files later in the week.

3. **Incremental backups** – all primary storage data that has changed since the last incremental backup are written to a set of tapes. With an incremental backup, since only changed files are backed up, the time and burden on server and network resources is kept to a minimum. The trade off of using a weekend full / daily incremental approach is that if system or data loss occurs mid-week, the full backup has to be restored first, and then each incremental backup has to be restored in successive order. If a failure occurred on Thursday, first the data from the last full weekend backup would be restored, followed by sequential restores from Monday, Tuesday and Wednesday night's incremental backups, which can significantly extend the time it takes to restore data.

For each type of data in an organization, different backup policies may be in effect. Here are some examples:

- Email data store and database backups may be performed as full backups daily to reduce the time it takes to recover from a lost email data store or database. Performing a restore from a series of full, incremental and/or differential tapes can take much longer than just a restore from a recent full tape.
- Most data is protected with full backups on the weekend and incremental backups during the week. Some organizations replace midweek incremental backups with midweek differential backups to reduce the number of tapes that would be used during a restore from later in the week.

Weekend and nightly backups are used to restore files that have been accidentally deleted or have become corrupted. Most organizations maintain between four weeks and 26 weeks of full backups and the last one to four weeks of incremental / differential backup data on site. This allows users to request the restoration of data from an earlier copy.

Historical data in the form of weekly and nightly backups may be required by financial staff when third party auditors request financial information as of a certain date. Lawyers, in response to legal inquiry or subpoena may be required to produce these files from historical backup tapes.

Weekend and nightly backups typically require a duplicate tape be created to keep a complete data set of all primary data off site in case of a site disaster such as pipes bursting, natural flooding, hurricanes, tornadoes, fires and other disasters.

A backup server is a standard server that has backup software loaded onto it. Backup servers are connected to a local disk subsystem and/or one or more tape drives within tape libraries. Backup servers also communicate with servers, desktops and laptops to be backed up by having backup client (agent) software loaded onto each system. The backup agent moves all data with that server over the network to a backup server. Since the backup process consumes client machine processing and storage cycles, backups are typically performed at night to minimize the performance impact on users and applications.

## Disk-based Storage for Weekend and Nightly Backups

Due to the complexity of managing tens to hundreds of tapes created by weekend and nightly backups, and the challenges of quick and reliable data restores from these tapes, most companies are moving to disk-based backup at their primary site while maintaining tape for offsite data protection. This combination provides faster, more reliable and more manageable local disk-based backups and restores of data while keeping the cost of storing data offsite on tapes at a minimum.

## Intraday Data Protection

Intraday data protection allows data to be protected multiple times per day. The copy is often made to the same disk subsystem or to a separate disk subsystem depending upon the technology. Intraday data protection allows data that is lost near the end of the day to be restored from a point in time more recent than the previous night's backup tapes. For example, if a SQL database is corrupted at 3:30 p.m. and an intraday copy was created at 12:00 a.m., then only 3.5 hours of data would be lost. Without intraday data protection, restores would be performed using the data from the last nightly backup, which for this example, could represent as much as 7.5 hours of lost data assuming the work day began at 8:00 a.m.

In addition, since restores are typically from a disk storage system, not tape, recovery time is measured in minutes, not tens of minutes to hours with tape.

There are 3 basic types of intraday data protection:

- **Snapshots** are available on certain disk storage subsystems and filesystems. Snapshot software either makes or retains historical versions of files on a frequency that is determined by the user. Snapshots can be scheduled as often as once an hour. Snapshots are limited to a fixed maximum number, typically 64 or 256, depending on the implementation. Once this maximum number is reached, the oldest snapshot is eliminated when a new snapshot is taken. Snapshots provide for intraday data restores, assuming that the deletion or corruption happened within the period in which a snapshot is available. If the deletion or corruption happened earlier than the oldest snapshot, a restore would have to be performed from traditional backup tapes. Also, if the primary disk system fails, not only is all of the primary data lost, but so are all of the snapshots. In both of these cases, traditional nightly backups would have to be used to restore the lost or corrupted data.
- **Application dumps** are available for certain applications, such as Microsoft SQL Server and Oracle (Recovery Manager – RMAN), which have a backup utility built into them. With application dumps, point in time database images can be created multiple times during the day to reduce the exposure to losing more than a few hours of data in the event of system failure or database corruption.
- **Continuous Data Protection (CDP)** is fairly new and not widely deployed at this time. Think of CDP as continuous snapshots that can also be replicated to another system, resulting in a historical record of data that is up to date anywhere from within seconds to

an hour. With CDP, if a disk subsystem fails or becomes corrupted at 12:25 p.m., a restore could be performed from the CDP data just seconds to minutes earlier resulting in minimal data loss. As with all technology, there is a trade off - CDP creates a byte-level replicated copy and therefore takes additional system, network and storage resources. As with other forms of intraday data protection, data that is protected by CDP must also be backed up nightly with traditional backup applications.

Snapshots, application dumps and CDP are not replacements for weekend and nightly backups. In each case, their function is to keep a minimum number of intraday copies to reduce the amount of data lost in the event of midday data loss or corruption. In addition, weekend and nightly backups are still required for long term history.

## Remote Office Backups

Remote office backups are usually performed by having someone at each remote office be responsible for backing up data in that office to a local tape drive. The challenge is that many remote office backups are not being performed correctly or reliably since these offices have no IT staff that can respond to the many issues that crop up when backups fail.

As a result, many companies are looking for new solutions to the remote office backup problem.

One approach to resolving this remote office backup problem is to have all remote office data replicated to a central office or data center where trained IT staff can centrally manage all aspects of the backup process.

Replication systems are designed to keep files between two or more sites up to date with each other. As changes are made at one site, the changes are replicated to one or more additional sites. New files are replicated in their entirety when they are first created, but when files are modified over time, most replication products only transmit the bytes or blocks of data that changed from one version of the file to another. Since just the changes to files are replicated, the amount of wide-area network bandwidth consumed between sites is minimal. This makes replication a good solution for moving data from one or more remote offices to a primary office where it can be backed up as a part of the nighttime backup run by trained backup IT staff.

Replication software must be installed onto each server at a remote office and at a replication server at the central site. When a nightly or weekend backup job starts up at the central site with the replication server as a backup client, all of the replicated files from all remote offices are backed up through the backup server to disk or tape.

Replication alone does not replace backup. Any accidental file deletion or any file corruption that occurs at the remote office gets deleted or corrupted at the central site replication server. When this occurs, a restore from some form of backup data set must be performed. Therefore, weekend and nightly backups of replicated data are required to maintain the history of these replicated files.

Since replication products do not consume lots of WAN bandwidth, many service providers offer remote site replication services to maintain updated copies centrally at their data centers. Replication software moves data from one or more customer servers across the WAN to the service provider's data center where it is centrally backed up.

The WAN traffic between sites during peak work hours is typically consumed for most companies, even without replicating remote office data to a central facility. For this reason, replication software can be configured to perform data replication to a central site only during off-peak times. In this case, backups should be configured to start when all data from remote offices has been replicated to the central facility's replication server.

## Offsite Data Protection

Most organizations keep many months worth of backup data on site to in order to respond to requests to restore files from users and to recover from storage and system failures. Months to years of backup data is also retained on tens to hundreds of tapes at an offsite environmentally controlled, secure, protected, tape storage facility. These can be recalled when onsite tapes are lost or found to be unreadable. One or more tapes may also have to be recalled in response to requests for archived project data, financial or legal data, or to recover data to new systems in the event of a system or site disaster.

It is not uncommon to see organizations keep a year's worth of weekly backups, many years worth of monthly backups, and up to seven years of yearly backup tapes stored offsite.

Organizations that are moving to disk-based backup systems at the primary site are also moving to offsite disk-based systems as well to dramatically reduce recovery time in the case of a site disaster and to eliminate tape media, drives and libraries, tape handling and tape storage service costs.

## Disaster Recovery

Many businesses are regulated and are required to provide active site disaster recovery systems and facilities. For each business application, there is a defined Recovery Time Objective (RTO) which specifies the maximum amount of downtime that the application will incur in the event of a site disaster. In addition to RTO, each application will have a Recovery Point Objective (RPO), which relates to the freshness of data after a recovery. For an application to be down for only one second (RTO) and to lose just one second's worth of data after a complete site disaster, the most costly and sophisticated replication, networking and clustering systems must be employed. There are many IT solutions that are much more cost effective and provide reasonable RTO/RPO (less than one hour) that would satisfy a great majority of a business' applications.

Disaster recovery service providers can offer a complete set up of servers, networks and storage, ready to go, to recover from site disaster. Unfortunately, the cost of these services is out of the reach of most mid-sized companies.

Some companies keep their own hot spare systems (servers, storage and networks) at a second location for their most business critical applications. Disk mirroring and replication systems can replicate data from a primary site, across the metropolitan or wide area network to a second site.

The most commonly deployed site disaster recovery solution employed involves rolling out cold spare storage subsystems, networks and servers. Most organizations have days to recover from a site disaster and will either keep spare systems at a second location or have a service contract with a provider who will provide spare systems within four hours or less from event notification.

Recovery data for these spare systems typically comes from backup tapes. Many companies are moving to disk-based backup repositories to reduce the recovery time after a disaster to minutes instead of hours to days. Also, a disk-based backup repositories can replace offsite tapes and the management overhead of offsite tapes.

## About ExaGrid

ExaGrid® offers a turnkey appliance that works in conjunction with your existing backup applications and is 30 percent the cost of standard SATA storage. ExaGrid provides data reduction through byte-level data de-duplication technology, which stores only changes for each version instead of storing full file copies. This unique approach reduces the amount of disk space needed to at least 20 to 1, resulting in a significant cost savings over standard SATA storage.

ExaGrid Systems, Inc.

2000 West Park Drive  
Westboro, MA 01581

**1 800.868.6985**  
**[www.exagrid.com](http://www.exagrid.com)**

