Choosing the Right Backup Recovery Solution

nsk inc

we make IT work

Information Technology Experts • (617) 303-0480 • www.nskinc.com

Table of Contents

Table of Contents 2
Abstract 3
Why Backup? 3
Considerations 4
RAID 5
RAID 0 5
RAID 1 6
RAID 5 6
Tape Backup 7
Pros 7
Cons 7
Should you use it? 8
Snapshot Backup 9
Pros 10
Cons 10
Should you use it? 10
Mirroring 11
Pros 12
Cons 12
Should you use it? 12
Data Replication 13
Pros 14
Cons 14
Should you use it? 14
Continuous Data Protection (CDP) 15
Pros 15
Cons 16
Should you use it? 16
Data Deduplication 17
Pros 18
Cons 18
Should you use it? 18
Conclusion 19
About NSK 19
Sources 20

nsk inc
2 Liberty Square, 7th Floor
Boston, MA 02109
617.303.0480 617.303.0481
www.nskinc.com @nskinc

Abstract

Small and medium businesses (SMBs) need to keep their businesses up and running just as much as larger companies. Yet they are challenged by several limitations when seeking out a solution that meets their mission-critical data recovery, application uptime, and data retention needs.

These include:
• Limited funds
• Limited IT resources
• Limited/inadequate solution options designed for SMBs

With these challenges in mind, this paper identifies and explores different types of backup recovery solutions and their benefits for SMBs.

Why Backup?

Think about the number of Word, Excel, Powerpoint, Outlook, and database documents you have created, edited and saved today. Now multiply that by a week, a month, a year. Then add in all of your colleagues' changes to this type of data. That is how much information is floating around in most company computers. Without an adequate data recovery system in place, a disaster could delete essential company and client documents that can never be replaced.

Studies have shown that:

- By the 6th day of a major data loss, companies experience a 25% loss in daily revenue. By day 25 it is 40%.[1]
- 43% of companies that experience a severe data loss disaster, and have no recovery plan, never re-open.[2]
- The cost of data breach can cost upwards of $204,000 per instance.[3]

Causes of data loss include:

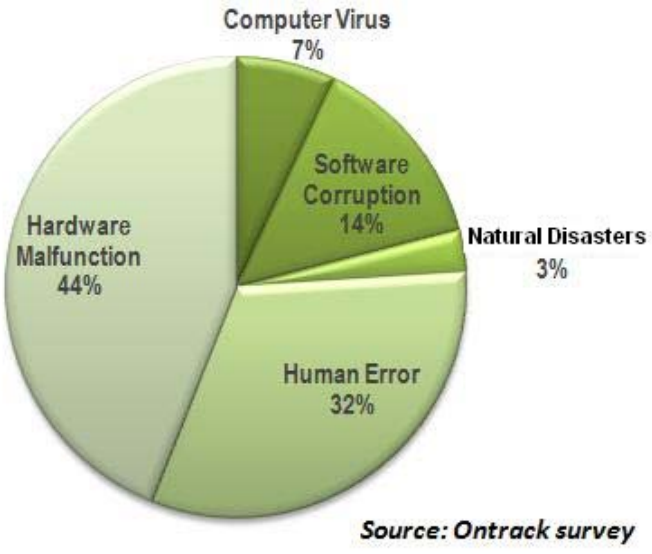

[1] Sequredata.com
[2] Sequredata.com
[3] Data Break Study, April 2010 Ponemon Institute

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

Data loss is preventable with the right backup recovery solution, often with little or no loss in productivity. Be sure to choose the right backup solution based on your needs.

Considerations

Backup recovery solutions exist for all of the different needs of SMBs. In order to determine which solution is best it is important to consider the following[4]:

- **How important is the data on your systems?** If crucial data lives on all of your machines, you need a backup system that extends on several servers and for several backup periods. This will ensure that no critical data is compromised in the case of a disaster.
- **How often does the data change?** If you update important files on a daily basis, more frequent backups are necessary.
- **How quickly do you need to recover the data?** Determine how much time you can allot for recovery after a disaster so you can choose the best solution.
- **What equipment do you need?** Backup hardware and software is often required in order to perform backups. Think about whether your budget allows you to purchase tape drives, optical drives, and/or removable disk drives.
- **Do you need to store backups off-site?** Yes you do! In order to save your files from a natural disaster you must store your backups off-site.

Data storage devices range from SAN, to VTL, and recently SSD (solid-state drive), which uses solid-state memory to provide the same sort of backup as a hard disk drive. Hard disk drives are electromechanical devices that contain spinning disks and movable parts, but SSDs use microchips and no moving parts. They are harder to break, quieter and update data at a higher speed than physical drives. Hard drives have replaced most tape backup systems. More and more tape backup bugs are fixed every year, which poses the question: Can tape backup stay relevant in the digital landscape?

After taking a look at your company and assessing its needs you can move forward with selecting the backup recovery solution that works best.

[4] Microsoft

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

RAID

It is important for your backup disaster services and your servers to be configured with Raid. We will discuss the different types to keep in mind when thinking about how to setup and configure your disaster recovery solution.
RAID (Redundant Array of Independent Disks) was first introduced in 1987 by David A. Patterson, Garth A. Gibson and Randy Katz at the University of California. A well known and widely used form of data protection setup, RAID is most commonly known as a method of dividing and replicating data among multiple physical drives. Despite being on separate drives, RAID data can still be access by one operating system.
Glossary:

- **Mirror-** Mirroring is when data is duplicated and written onto an identical disk. This creates redundancy and will help prevent data loss if disk failure is to occur. However, with this method, associated RAID levels cannot be entirely efficient or cost effective due to the secondary drive reserved for duplication.
- **Striping-** "The amount of data that each drive writes and reads in an array at one time."[5] Essentially, it is how many (almost) simultaneous writes it can do per disk.
- **Parity-** Used for fault tolerance. It calculates the data in two drives and stores it on a third.
- **Rebuild-** Rebuilding is when data is recovered when a disk failure occurs.
- **Single/Multiple bit corruption-** These forms of corruption indicate that an error occurred during the writing process.

There are nine common forms of RAID; however, there are many more that have been developed over the years. The different types of RAID are named by the word RAID followed by a number (ex: RAID 0, RAID 1, etc).

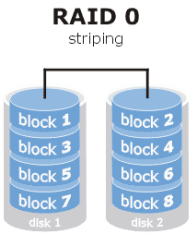

RAID 0

This form of RAID uses a striped disk array, which means data is separated into blocks which are written onto different disk drives. Also, it performs best when data is striped across multiple drives[6].

[5] www.buzzle.com/articles/raid-levels-explained.html

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

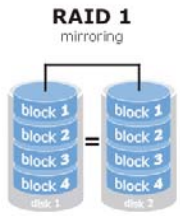

RAID 1

RAID 1 introduces mirroring—one hard disk stores information and the other stores duplicated data. Due to the duplication, data redundancy is extremely high in this level, which leads to a higher cost of implementation. However, in the case of disk failure, no rebuild is necessary since all of the original data has already been duplicated

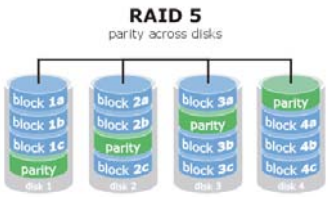

RAID 5

This level of RAID uses block level striping, but there is no hard drive dedicated for storing parity data. Instead, parity data is distributed across all hard disks. Information is also backed up by using the parity data which accumulates to the size of one hard disk. This level of RAID is one of the more popular versions and has a low cost of redundancy.

[6] JetStor

Tape Backup

In a world that is becoming increasingly digital, RAID employs a physical disk drive, which is why it is important to note where physical data storage originated—tape.

In the event of a hard disk crash, or other catastrophic computer failure, tape backup ensures that data will not be lost. It involves copying data, either manually or with appropriate software, onto an actual tape cartridge. Tape backup also includes the ability to restore data that has been backed up back to hard disk storage devices.

1994 marked the beginning of fibre-channel (FC) adapters, when developers began creating hardware to service high-speed data transfer. At the turn of the twenty-first century, 3 out of the top 10 fastest growing profitable US storage companies were tape companies. Tape backup became the most economical way to store large quantities of data.

In 2003, virtual tape libraries (VTL) entered the market, allowing for faster access to the backup drive for recovery. Parallel ATA (PATA) and Serial ATA (SATA) disk arrays are the most cost effective primary storage components. The use of disk arrays paved the way for an increase in storage capacity. VTL also eliminated streaming problems associated with uploading a mass of data. It should be noted that VTL is not really tape. It is a concept originally created for IBM mainframes, but it is now exploding in the open systems arena. **VTL** is physically a highly intelligent optimized disk-basedstorage appliance. Because a **VTL** completely emulates a standard library, the introduction of **virtual tape** is seamless and transparent to existing tape backup/recovery applications. [7].

Pros

- Information on a tape backup is stored on a tangible, physical tape, which makes it dependable and long-lasting. Tape backups can be stored for decades without any damage to the content. In fact, many tape storage devices that were made over 20 years ago are still readable on new drives.
- Tapes are portable and can easily be taken offsite.
- Tape backup is well established and requires low heat and power costs to maintain.

Cons

- Backing up to physical tapes involves spending a lot of time and money because data rate is slow due to the tape's density. Additionally, tapes need to be manually inserted and rewound, which requires hours of employee dedication.

[7] http://www.webopedia.com/DidYouKnow/Computer_Science/2006/virtual_tape.asp

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

- The cost of tape backup, per GB, is more than the cost of digital backup and the speeds of restoring from tapes are slower.
- If tapes are reused over and over again the quality of the data could deteriorate. Additionally, if not handled with care, a tape could be damaged.
- Restoring a specific file can be difficult because it is necessary to go through an entire tape to find the data. There is no random access as with hard drives.
- The drives that tapes backup are sometimes larger than the storage capacity of a tape, so multiple tapes need to be used. This can require extra time from employees or the purchase of a tape library.

Should you use it?

After reviewing the pros and cons of tape backup, it is evident that tapes are not going to fall out of favor anytime soon. Despite the cost, they are a stable and durable and should not be passed over for digital options. It all depends on the amount of data that needs to be stored and how quickly it needs to be accessed. If a business has over 2 TB of data, digital options may be cheaper and data will be easier to retrieve. For sensitive data or to make sure that files are never deleted, tape backup is more reliable. The best option for backup is to use a hybrid of the two options—save an entire server onto a digital drive, but copy extremely important information to a tape as well.

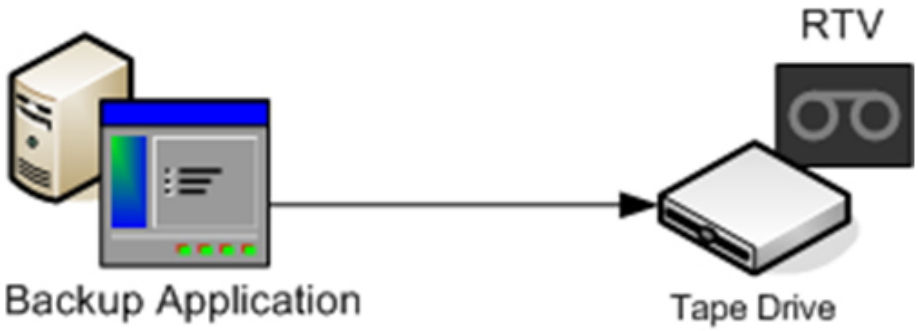

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

Snapshot Backup

Snapshot storage backup recovery works almost exactly how it sounds: by taking a photo of the computer hard drive at a particular moment in time. The replication occurs by copying a disk block when it is written or by using a split-mirror set of disks. When it is time to restore from snapshot backup, the result is almost instantaneous. Snapshot backup can deliver full, partial, file and even some differential database backups.

Unlike hard disk or tape backup systems, snapshot backup is not a physical backup drive, but a method of backing up data. It takes all of the contents of a computer drive and makes a copy, places the copy on another disk and modifies the copied file when changes are made. An example of snapshot backup can be found on the Mac Operating System in the form of the Time Machine. Time Machine takes incremental snapshots that allow a user to restore an entire system, multiple files, or a single file in minutes.

Two main types of snapshot backup exist: copy-on-write (AKA low-capacity) and split-mirror. Copy-on-write creates a new snapshot every time data is updated or new data is entered. This is helpful in the case of program malfunctions or if a file becomes corrupt. To restore all of the data on a network, however, all past snapshots must be accessible. Unlike copy-on-write, split-mirror backup creates a snapshot of the entire volume of data every time it runs. This option makes the data available offline, which simplifies the process of recovering, duplicating, or archiving the data. The process is slower, however, and necessitates more storage space.

Copy on write illustration (from IBM)

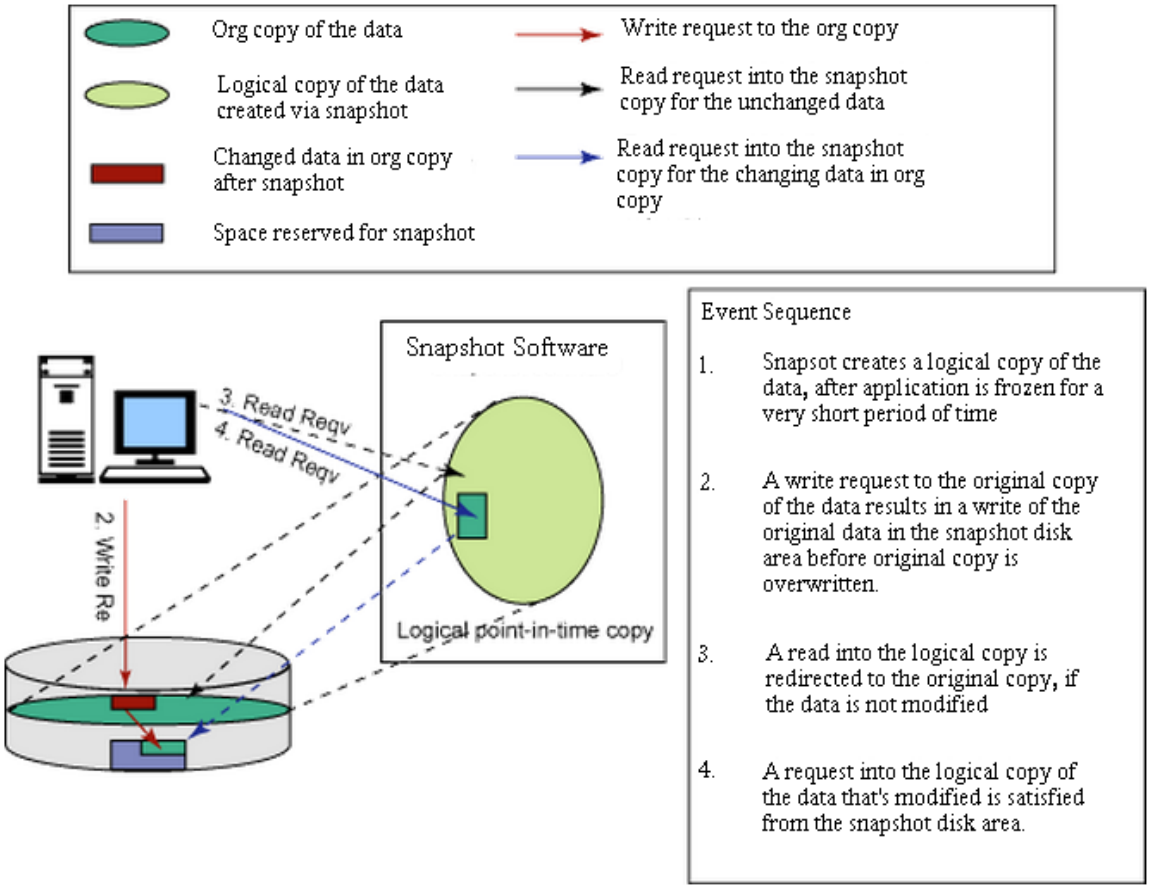

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

Pros

- The backups are small and can be created extremely quickly, with little effect on the server.
- It is possible to schedule snapshots and take as many as are necessary for system backup and recovery.
- Corrupted and/or deleted data can be restored using snapshots. Additionally, in case of user error, snapshot backup makes it easy to revert back to the snapshot that was taken before the error occurred.
- Restoring a server from backup can be accomplished instantaneously.
- Snapshot files are small
- Snapshots enable better application availability, faster recovery, easier back up management of large volumes of data, reduces exposure to data loss, virtual elimination of backup windows, and lowers total cost of ownership (TCO).[8]
- Because some snapshot-based systems keep the data in its native format, it is possible to simply switch to the replicated copy without restoring an entire system.

Cons

- Another backup of the entire server must be made to make up for the possibility that the server itself may become corrupt. This means that with snapshot backup, onsite and offsite backup needs to be used in order to guarantee absolute data recovery.
- It is possible for the snapshot to run out of disk space, in which case no other backups can be made.
- Too many snapshots slow down the production volume of a computer. However, this only occurs with copy-on-write snapshot data. With other snapshot systems, such as redirect-on-write, the lag in production volume is hardly noticeable.

Should you use it?

Despite its minimal shortcomings, snapshot backup advantages seem to greatly outweigh the negatives, making it an extremely practical backup solution. Snapshot backup technology makes it easy to note changes made to computer data in order to recover a file if the need arises. It is not only beneficial for large companies that store masses of data, but also for personal computer backup as well.

[8] ibm.com

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

Mirroring

So far we have discussed tape backup and snapshot backup as examples of data recovery solutions. Each of these methods copy data on a computer in case disaster recovery is necessary. Mirroring is another type of data recovery solution that is often used to create multiple drives with the same information. It works by linking two systems together that will continually match each other. For example, when "File A" is created on a computer, it will be saved on two disks at once. Mirroring is used for a variety of reasons—for data backup, for faster downloads in different geographic locations, and to balance a server load, also known as clustering.

Mirroring can occur in two regards: on disks, or on server. RAID (Redundant Array of Independent Disks), is an example of mirroring that is stored on two separate physical disks in order to ensure that the data is always available. Mirrored data can also be stored on two different servers in two separate locations. This is done so that the same information is quickly accessed at various sites. Mirroring is also part of cloud technology.

If a computer malfunctions and disaster recovery must be initiated, mirrored solutions work through a hot swap system or a hot standby disk system in order to save data. In a hot swap system, the system itself signals a disk failure and switches over to the mirrored disk. Users frequently cannot even tell that this was done, except for an alert that may pop up on their screen. Hot standby is a more sophisticated system during which a backup occurs from the active disk onto the remaining disk. Then, a new disk is installed creating a new backup system, while the malfunctioning disk is discarded. During this process, little data is lost and minimum disruption occurs. With mirroring, data recovery can be performed swiftly and rapidly.

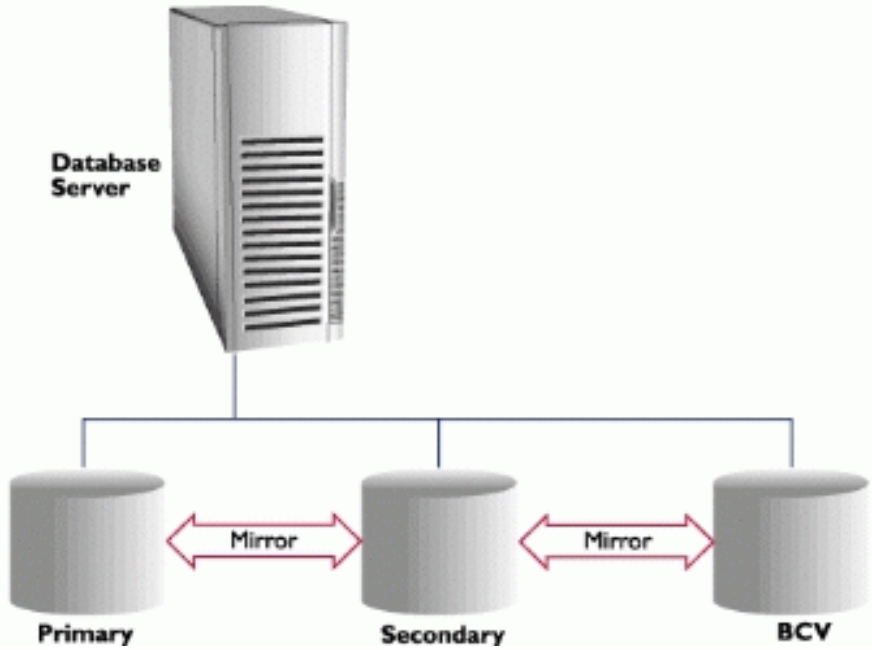

Pros

- Mirroring is robust and efficient because the changes synchronize without any lag or data loss.
- Mirroring can require less bandwidth than other methods of replicating data. Also, replication is not limited by geography so servers located far away can mirror each other, provided there is enough bandwidth.
- Mirroring allows for real time data security because the data is written to two drives at the same time.
- Because data is always copied onto a second drive, backup can be performed quickly and easily without slowing down productivity.
- Mirroring delivers a high degree of protection against disk and hardware failure.

Cons

- The downside of real time data security is that if the user incorrectly changes a vital aspect of the Operating System, the backup will be affected as well.
- Because backups occur frequently during mirroring, the second drive receives a lot of wear. Also, mirroring adds some extra load on the principal server because its function is to keep the database mirror synchronized.
- Mirroring can lead to less storage space available for data, because both drives need to have enough space to support the entire hard drive.
- Having multiple disks with the same information may arouse a false sense of confidence that can lead to poor backup maintenance. Disk mirroring does not prevent against viruses, file deletion, and other logical loss of data.
- If mirroring is removed for some reason, the whole process of configuring mirroring has to be completely redone.

Should you use it?

If using it as a way to backup a single computer drive or server, mirroring can be quite successful. If the intention is to back up many computers, the process can get more expensive and time consuming. Mirroring is also a good option for storing sensitive data because a copy of the data always exists in an easily-accessible location. However, somebody looking for a backup recovery system without needing to document every instance of data change may want to choose a simpler backup recovery software.

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

Data Replication

Database replication involves creating and maintaining multiple copies of the same database for recovery and backup purposes. It requires transferring files from one computer or server to another so that all users share the same data. This results in a distributed database that gives users access to necessary information without interfering with others' work. Normalization, which eliminates ambiguous data and inconsistent files among users, serves as the primary motivation for database replication. A distributed database management system (DDBMS) guarantees that anything that is added, deleted, or adapted on one server is reflected automatically in all other locations, creating an overall consistent database.

Database replication can be performed in three different ways:

1. **Snapshot replication:** Data is copied from one server to another. (Read our blog post about snapshot recovery here.)
2. **Merging replication:** Data from multiple databases is combined into a single database.
3. **Transactional replication:** Users receive full initial copies of the database and then receive periodic updates as data changes.

The most common method of database replication is multi-master replication, which allows data to be updated by any user and stored on a group of computers. The alternative to this process, multi-slave replication, gives sole responsibility to a master computer that makes changes whenever a user updates data. Multi-master replication allows for a quicker and more efficient way to replicate data. Replication occurs through database transaction capturing, which obtains changes made to the database synchronously, asynchronously, or both, depending on the method of capture.

All of the various types and methods of database replication can make it complicated to implement, especially when data increases in size and magnitude.

Pros

- Database replication allows all system users to easily and efficiently share data. It recognizes data modifications and copies the information so that it can be accessed.
- During multi-master replication if one master fails, another one can continue to update.
- Multi-master replication allows masters to be location across several physical sites and distributed across an entire network.
- Multi-master replication is a very powerful and flexible.
- It is possible to create multiple replication groups within a single database.

Cons

- Database replication turns out to be complicated when data increases in size and magnitude.
- Conflict resolution can become increasingly complex due to the number of nodes involved in multi-master replication.
- Asynchronous replication, associated with multi-master replication, increases the possibility for inconsistent data changes, or "conflicts" to occur during database transaction capturing.
- Setting push jobs using asynchronous replication to run at very short intervals can use up an excessive amount of system resources.
- Usually only used with database based systems and is not applicable to an entire site.

Should you use it?

Data replication serves as an incredibly powerful and functional solution to data backup systems. However, it is not an easily managed solution, and comes with a great deal of complexity. It is recommended that multi-master replication be used to protect a critical database in case the primary site becomes unavailable due to system or network outages. For applications that require multiple access points for users, data replication can help distribute a heavy load and provide continuous availability for located data. Alternative backup needs can be better accommodated by other backup systems.

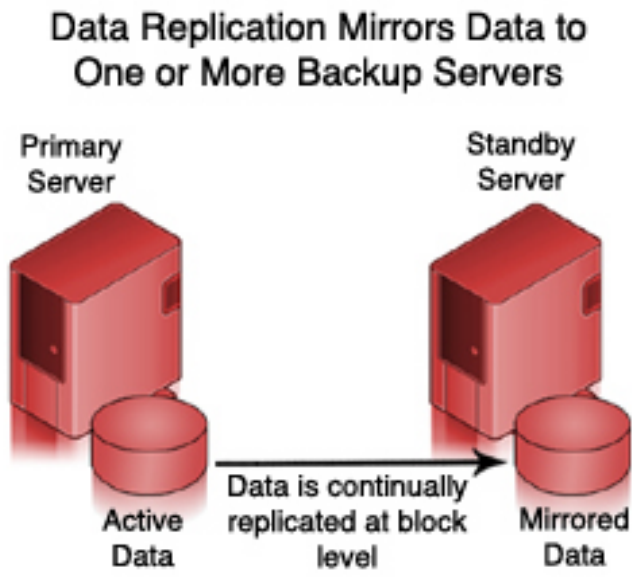

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

Continuous Data Protection (CDP)

Continuous data protection (CDP) differs from traditional backup recovery software in one essential way: whereas backup methods require data to be copied to another entity, such as a tape or a disk, CDP data is copied and then stored on the local storage area network (SAN) or the network-attached storage (NAS) system. CDP works by making frequent, incremental copies of the data itself.

Essentially, CDP tracks every change made and captures continuous changes to data. CDP serves as a baseline reference to the original state of the data, it tracks the state of a file when a change occurs and tracks the change in the backup system, and it allows granular recovery for multiple point-in-time states of the data.[9] When an incident occurs and backup recovery is necessary, CDP allows the user to go back and reset a file, dataset, or entire database to the state it was in before.

Pros

- Using CDP after the initial setup does not require intensive IT training and the backed up files can be easily accessed by the average desktop and laptop user.
- CDP integrates into established data protection frameworks very easily and can protect data on other storage tiers.
- CDP preserves a record of every change made to a computer.
- If a system becomes infected with a virus or a while is corrupted and the problem is not discovered until later, it is still possible to recover the most recent clean copy of the affected file.[10]
- Data recovery is possible in a manner of seconds, much less time than with other backup systems. However, this is contingent on file size and network speed.
- Installing CDP hardware does not put existing data at risk. Furthermore, there is no gap in data when a restore occurs.

[9] Search SMB Storage
[10] Channelpro SMB

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

Cons

- The hardware required to set up CDP can be costly and complex to install initially, and low-cost alternatives do not perform as well.
- There is a lag in computer performance when data is being replicated because it must be written to the target disk (SAN or NAS system).
- Data recovery is dependent on network speed and file size, so it may not always occur quickly, especially if a file was frequently updated.
- CDP backup has a varying degree of granularity, which is tied to the last time a file was created or saved. Therefore, it may not always catalogue every single change that occurs.
- "System performance (CPU, network, etc.) may be adversely impacted by the protection process, especially when there are limited windows of opportunity to replicate data to corporate storage systems."[11]
- If a file is mistakenly restored from CDP backup, it can be difficult to undo.

Should you use it?

Using CDP as the only means of backup can be expensive, especially if a small business has less than 1TB of data to store. However, for businesses with a lot of data, even if it is in varying degrees of importance, CDP may be the perfect data recovery solution. CDP reduces the impact of computer failure or data loss and can make a significant difference for businesses that depend on critical files. Small businesses, in particular, benefit from CDP because it does not take an entire IT department to run.

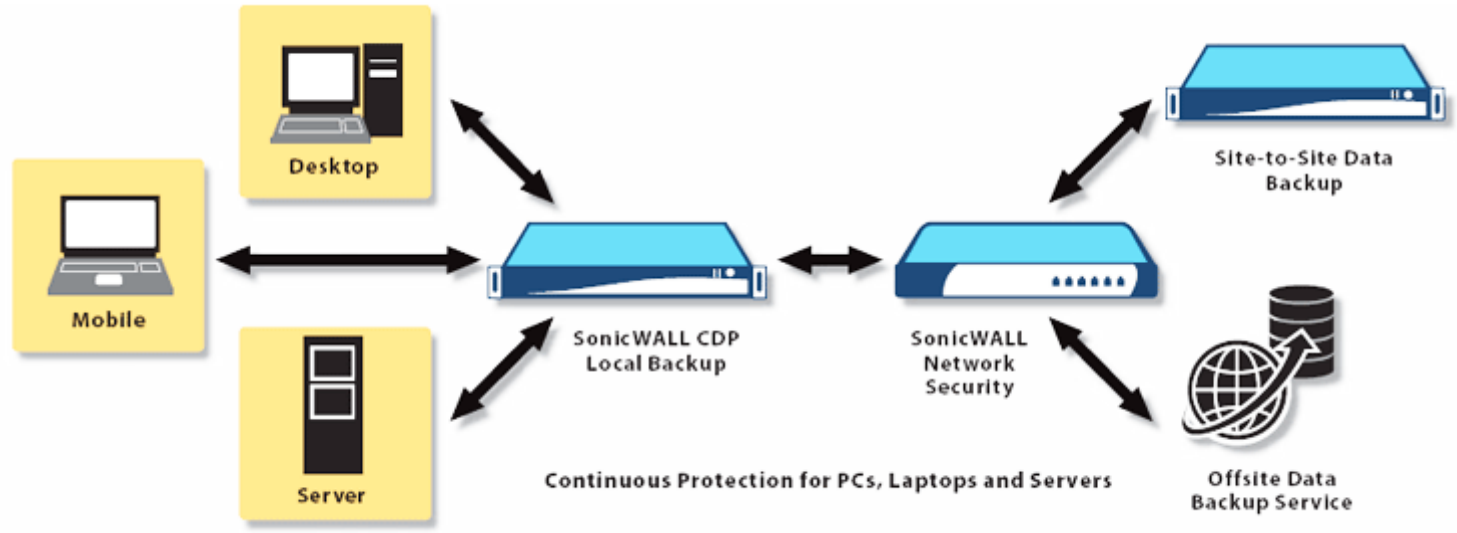

[11] Search SMB Storage

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

Data Deduplication

Data deduplication, although not traditionally considered backup software, can be quite handy when backing up large instances of data. The deduplication process works by identifying unique chunks of data, removing redundant data, and making data easier to store. For example, if a marketing director sends out a 10MB PowerPoint document to everyone in a company, and each of those people saves the document to their hard drive, the presentation will take up a collective 5G of storage on the backup disk, tape, server, etc. With data deduplication, however, only one instance of the document is actually saved, reducing the 5G of storage to just 10MB. When the document needs to be accessed the computer pulls the one copy that was initially saved.

Deduplication drastically reduces the amount of storage space needed to back up a server/system because the process is more granular than other compression systems. Instead of looking through entire files to determine if they are the same, deduplication segments data into blocks and looks for repetition. Redundant files are removed from the backup and more data can be stored.

There are three ways for deduplication to occur:

1. With **inline deduplication**, data is deduplicated before being stored. This does not require any additional space to store the data prior to deduplication.[12]
2. **Post-process deduplication** briefly places all of the backup data on a disk-based staging storage prior to being deduplicated.[13]Then, the data undergoes the deduplication process. Although this method requires more space, it enables faster backups and recovery.
3. "**Source-side deduplication** typically uses a client-located deduplication engine that will check for duplicates against a centrally-located deduplication index, typically located on the backup server or media server."[14]

Despite any of these methods being used, data deduplication is not a stand-alone product and must work in conjunction with other storage solutions.

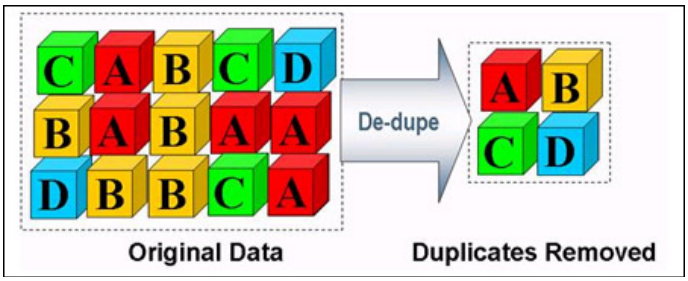

[12] Continuity Central
[13] Continuity Central
[14] Continuity Central

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

Pros

- Compressing data gives SMBs more bang for their buck, because they can make the space on their current storage devices go further by removing duplicate data.
- "Less data can be backed up faster, resulting in smaller backup windows, smaller (more recent) recovery point objectives (RPOs) and faster recovery time objectives (RTOs)."[15]
- Data deduplication speeds up backup, replication and the disaster recovery processes.[16]
- Deduplication can lead to sizeable savings in terms of time, resources and budget.
- Because deduplication decreases file size, it helps remove the amount of media required to provide a SMB with good quality data recovery.

Cons

- There is a small potential of data loss when data is deduplicated because a deduplication system stores data differently than how it is written. Therefore, the reliability of the data depends on the deduplication system. However, the development of technology over the years has decreased the chance of data loss.
- If using the inline deduplication method, data that does not deduplicate well has the possibility of being erased.
- The source side deduplication method can become easily overloaded with large files, which can slow down backups.
- Data deduplication is not a stand-alone product, it must be used with other backup software.
- Some deduplication methods, such as post process, require more difficult configurations to make them function correctly.

Should you use it?

With the abundance of data deduplication tools on the market, most of the cons associated with deduplication can be avoided by choosing the correct software. For SMBs trying to save money on a backup solution, deduplication can significantly decrease the amount of necessary storage space. With deduplication, backup and recovery can be attained much more quickly and without redundancy.

[15] Search Storage
[16] Search Storage

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc

Conclusion

The most important thing to remember when choosing data recovery software for your SMB is to make sure your solution is tailored to your company needs. Keep in mind how much data you have, how important it is, and the resources you have for making sure your data is backed up. You can never go wrong with backing up too much or too often as important documents tend to be the lifeline of a company.

Consider turning to an experienced IT company that can support your most critical data on highly secure systems that are well protected against any type of damage. With a complete backup solution system from NSK Inc there is no risk of ever losing your data. Pavis, for example, is a system developed by NSK with the help of our partners. Pavis combines the pros explained in this document into one unique system. It is a hybrid system that combines RAID, Snapshot, CDP and offsite storage to the SMB as a robust backup, business continuity and disaster recovery system that is comparable with the large enterprise systems but does not cost the SMB the large enterprise price.

If you are having difficulty determining your needs or figuring which solution is right for you, sign up for a free onsite consultation to help you with your backup and recovery problems.

About NSK

NSK Inc is a leader in information technology consulting, with a focus on IT management for SMB companies. Headquartered in Boston, MA with an additional office in Palo Alto, CA, the company offers a wide array of IT services for business driven information challenges. They provide service and support for small and medium-sized businesses and groups working within large organizations. For more information, please visit http://www.nskinc.com.

2 Liberty Square, 7th Floor
Boston, MA 02109
617.303.0480 617.303.0481
www.nskinc.com @nskinc

Sources

Raid
www.buzzle.com/articles/raid-levels-explained.html
www.acnc.com/raid

Snapshot
http://sql-articles.com/articles/dba/pros-and-cons-of-database-snapshot/
http://www.simple-talk.com/sql/database-administration/sql-server-2005-snapshots/
http://msdn.microsoft.com/en-us/library/ms189548.aspx

Mirroring
http://searchstorage.techtarget.com/answer/RAID-types-and-benefits-explained
http://bettereducation.com.au/it/yaf_postsm616_Pros-and-Cons-for-Database-Mirroring-By-Shailesh-patangay.aspx
http://www.backupcritic.com/articles/raid-for-home.html
http://www.mssqltips.com/sqlservertip/1647/evaluating-sql-server-database-mirroring-as-a-disaster-recovery-solution/

CDP
http://thebackupblog.typepad.com/thebackupblog/2008/09/cdp-vs-backup.html
http://inmage.dcig.com/2008/07/cdp-and-backup-software-the-fu.html
http://searchsmbstorage.techtarget.com/tip/Continuous-data-protection-Pros-and-cons-for-SMBs
http://www.channelprosmb.com/article/20721/Pros-and-Cons-of-Continuous-Data-Protection/

Data Deduplication
http://aim.uoregon.edu/research/ebriefings/pdfs/eb-klearman.pdf
http://searchstorage.techtarget.com/feature/How-data-deduplication-works
http://searchdatabackup.techtarget.com/podcast/The-pros-and-cons-of-data-deduplication-tools

Database replication
http://searchsqlserver.techtarget.com/definition/database-replication
http://www.tech-faq.com/database-replication.html
http://download.oracle.com/docs/cd/B12037_01/server.101/b10732/repmaste.htm
http://www.dbspecialists.com/files/presentations/mm_replication.html

Images
Cover: http://i-data-recovery.com/
Page 3: http://intelligentpc.net/blog/?p=10
Page 4: http://www.vwsolutions.com/VisionBackup/
Page 5: http://www.mnmsystems.com/qnapts4108tb-1.aspx
Page 6: http://www.hdd-tool.com/raid/raid-0-features-advantage-and-disadvantage.htm
Page 7: http://www.prepressure.com/library/technology/raid
Page 8: http://blog.everycity.co.uk/2008/10/07/making-servers-faster-raid-1-vs-raid-5-performance-differences/
Page 9: http://insnewsletter.com/newsletter.asp?article=9
Page 10: http://www.emaglink.com/virtual-storage-manager-recovery.aspx
Page 12: http://www.tkshare.com/drive-snapshot-v14014979.html
Page 13: http://technet.microsoft.com/en-us/library/cc966458.aspx
Page 14: http://www.digitizd.com/2009/08/11/two-computers-powerful-and-portable-one-low-price-through-aug-12/
Page 15: http://www.cascadesolutionsllc.com/uploads/images/servers.jpg
Page 16: http://www.openminds.co.uk/windows-solutions/data-replication/data-replicationmore-information.html
Page 17: http://www.ibackup.com/professional/online-backup-features.htm
Page 18: http://www.sonicguard.com/CDP4440i.asp
Page 19: http://www.eweek.com/c/a/Data-Storage/How-to-Leverage-Data-Deduplication-to-Green-Your-Data-Center/
Page 20: http://blog.appassure.com/blog/appassure-software/0/0/replay-data-deduplication-and-compression-software/

nsk inc
2 Liberty Square, 7th Floor
Boston, MA 02109
617.303.0480 617.303.0481
www.nskinc.com @nskinc

Page 21: http://www.affineontech.com/data-recovery-tools.htm

nsk inc 2 Liberty Square, 7th Floor Boston, MA 02109 617.303.0480 617.303.0481 www.nskinc.com @nskinc