

D2D & D2D2T.... The New Backup



**What do you mean
we lost our data?!**

Abstract

The fastest growing party of any IT installation is normally its lifeblood: its data. Data volumes have ballooned significantly in the 21st century, and the rate of annual growth is also increasing. Moreover, new national standards mandating how often and in what manner data is to be protected arrive daily.

This has forced changes in the traditional ways that data is backed up. Saving data onto tape or optical disk is becoming too slow for those environments where the volumes are enormous, or there is little time in the operations schedule to create backups, or both.

Disk-to-disk (D2D) and Disk-to-disk-to-tape (D2D2T) technologies have evolved as a way to meet the needs of IT installations that simply cannot continue to do their backups in the traditional manner. Although similar devices, this paper will treat D2D and D2D2T appliances as a group. The promise disk-based virtual tape devices hold for improving the data protection function needs to be understood.

A dedicated manufacturer of storage solutions, GST, Inc. produced this GST StorFacts™ Report as part of its commitment to provide information leading to the improved safeguarding of mission-critical applications and data within the IT industry.



Table of Contents

Abstract	1
Forward	3
Executive Summary	4
Background	5
Key Areas for Disk to Disk	6
1. Factors Relevant to D2D and D2D2T	6
2. What is a Disk to Disk Appliance?	7
3. What is a Disk to Disk to Tape Appliance?	8
4. What are the General Benefits?	9
5. What is the Hierarchical Storage Capability?	11
6. What types of Disk are used?	12
7. Optimizing Performance	13
Feedback	14
Related GST StorFacts™ Reports	14
About the authors	14
About GST, Inc.	14
Trademarks	14



StorFacts™ Report

D2D & D2D2T

Forward

This GST StorFacts™ Report is part of a series of StorFacts™ Reports that GST provides to aid in the consistent and effective safeguarding of mission-critical applications and data.

Other StorFacts™ Reports in the series include:

Top Ten Storage Guidelines for Storage Readiness

http://www.gstinc.com/white/compliance_10.pdf

Current Compliance Issues for Storage Professionals

<http://www.gstinc.com/white/compliance.pdf>

SLR and Your iSeries

<http://www.gstinc.com/white/slr.pdf>

AIT Tape Technology

<http://www.gstinc.com/white/ait.pdf>

The Need for Mirrored Backup

<http://www.gstinc.com/white/mirrored.pdf>

Backup Insurance, Are You Covered?

<http://www.gstinc.com/white/insurance.pdf>

GST StorFacts™ Reports (white papers) can be found on GST's web site at:

<http://www.gstinc.com/white/index.html>



Executive Summary

As part of GST's constant monitoring of developments in the storage, backup and data protection markets, GST talks to engineers, other storage manufacturers, our own distributors and buyers of data protection backup and data storage products. In doing so, we have talked to a number of IT installations and storage vendors that are using Disk-to-Disk (D2D) and Disk-to-disk-to-tape (D2D2T) appliances.

An alternative to traditional tape backup, D2D and D2D2T both are providing significant benefits in IT shops today. Although more costly than tape backup media, disk is so much faster that it deserves consideration. Disk-to-disk (D2D) backups using SATA inexpensive disks for the receiving end of the D2D backup operation have become a major new solution in the backup world. D2D does not eliminate the need for tape or other removable-media backup, since with D2D the backup data stays on an appliance in an IT center and the object of backup is to get the data away from the IT center where it can be safely stored. When D2D backups are combined with offloading to tape, it is called disk-to-disk-to-tape (D2D2T). D2D2T appliances have a backend Fibre connection that can be attached to a tape drive, autoloader or tape library for use in offloading the backup data from the appliance to tape for storing off premises.

D2D2T appliances provide a means for executing much faster backups that include offloading to tape later at a more convenient time. Another advantage of D2D2T backup technology is that it provides the ability to retrieve data directly from the secondary disk array (the 2nd D in D2D2T) at much faster speeds than tape offers.

Choices of disk and tape for D2D and D2D2T appliances usually provide a range of technologies. 2005 will see further announcements in both D2D and D2D2T backup appliances. At GST, D2D appliances and D2D2T appliances are field upgradeable with more and faster drives and additional specialized components.

Clearly, organizations are looking for ways to eliminate the drag on server resources by the backup/restore process. D2D initially gained a *magic bullet* aura as a cure-all for the problem of ever-shrinking windows of time for performing backups. However, the popularity of D2D as a backup solution does not eliminate the need for a removable backup media like tape to permit easy and economical removal of critical backup data to a safe remote location.



Disk-to-Disk Technology

The report that follows provides information on the promise of Disk-to-Disk technology in its two most popular forms: D2D and D2D2T. This information can be useful to organizations assessing the value of various disk technologies when planning backup infrastructure and acquiring new storage resources.

Background

Traditional backup to tape is being replaced in larger installations with what's being called *virtual tape*, where high-performance racks of inexpensive disk drives look like tape to the server. A Disk-to-Disk (D2D) unit writes backup data directly to disk. When the disk appliance also includes a unit to offload the data to tape cartridges, it is called a Disk-to-Disk-to-Tape (D2D2T) appliance. The result is backups performed at disk speeds rather than tape speeds. This permits servers to be backed up at server rates with backup data going directly to disk for online storage and access later. With a D2D2T appliance, backup data is then offloaded to tape for safe removal to a remote location. SATA hard drives are not removable and therefore offer no disaster recovery protection without the tape backup combination.

A typical 6- or 10-cartridge magazine generally holds high performance cartridges that hold many TB of backup data on traditional tape media. The media can be WORM (Write Once Read Many) if compliance regulations so stipulate that level of protection. Inherent nearline disk capacity is in the terabyte range.

In whatever form it takes, backup of programs and data that reside on disk is one of the most critical functions in any enterprise today. The Gartner Group, a leading IT industry analyst organization, has listed the following phenomena associated with a business interruption:

- Nearly 75 percent of all US businesses have experienced an interruption.
- 20 percent of all small to medium businesses suffer a disaster every 5 years.
- 43 percent of all US companies never re-open after an unexpected business interruption and 29 percent close within 3 years.

And perhaps most significantly,

- 93 per of all businesses that experience a significant data loss are out of business within 5 years.



Furthermore, a recent study by the Contingency Planning Group shows that the financial impact of one or more hours of downtime by a main server can be quite costly:

- Nearly half of all companies reported that each hour of downtime costs them at least \$50,000.
- One in four companies reported their cost of downtime ranges from \$250,000 to \$1million or more.

In a survey of IT managers taken by the Enterprise Storage Group:

- 58% had already implemented some form of D2D backup.
- 25% planned to purchase D2D backup solutions within the next 24 months.

Bottom line: while selection of the right disk technologies is important, the failure to adequately backup both programs and disk-based data files can threaten the organization's viability. Organizations that loose access to their critical computer files for periods of as little as one week can find themselves out of business within the next several years.

Implications: The decision on the type of disk technology to employ, although important, should be made within the context of the larger decision on what backup strategy to employ.

Key Areas for Disk to Disk

1. Factors relevant to D2D and D2D2T

There is no one backup strategy that is right for every situation. Each set of processing requirements and business conditions dictates a unique approach to data protection. To prepare an effective long-term data protection/backup strategy, it is important to understand your IT needs and business goals over the longer range. Key factors that must be considered are:

- **Capacity online** – How much data needs to be kept online?
- **Capacity nearline** – How much data needs to be accessible?
- **Scalability** – How rapidly are data volumes expected to increase?
- **Performance** – How fast does the stored data need to be processed?
- **Reliability** – What is an acceptable failure rate? What is an acceptable repair time?
- **Availability** – What is the acceptable time that critical applications can be down?
- **Backup** – What are the backup schedule needs?
- **Recovery** – Are recovery procedures in place?
- **Talent** – What in-house or contracted skills are needed for data protection?
- **Budget** – What is the budget for backup and recovery?



Gaining an understanding of these critical areas will lay the groundwork for determining the best mix of data protection and disk technologies over the next planning horizon. This will involve more than just IT personnel. Often, members from the Risk Assessment department, from Internal Audit, Budget Planning, Sales Forecasting and other business planning groups in the organization are needed to establish an understanding of these factors.

2. What is a Disk to Disk Appliance?

Disk to Disk (D2D) appliances are generally built around a server with enough disk to accept the backup data from another server. D2D technology permits backing up servers at disk-to-disk speeds. Later, backup data on the disks in the D2D appliance can be offloaded to tape media when production demands are less. Data on disk in the D2D Appliance can be accessed at high speeds for rapid recall of backup files.

GST's StorageDR™ D2D Backup Appliances are designed to deliver seamless integration of removable storage, SAN/NAS technology, and management software, targeted to meet the needs of Small-to-Medium-Business (SMB) and Enterprise customers for data replication, backup, archiving, and recovery.

Most D2D appliances use a range of configurable components to provide a variety of speeds and capacities. Variables that are packaged together into an appliance include multiple CPUs of varying CPU chip levels, ranges of cache memory, 10K and 15K RPM disk drives of varying capacities, Gigabit switches, SCSI and iSCSI connectivity and RAID offerings.

High performance backups and restores are accomplished using Serial-ATA disks that handle data at rates of half a TB an hour or more.

GST D2D Appliances utilize rack-mountable components and can be housed in GST's 11U, 25U, or 42U rack for easy installation and maintenance. A D2D starter configuration might contain a single processor and 2GB of memory, 2MB of Cache, 73 GB of 10K disks, single SAN switch. Additional processors up to three can be added later. Additional 1, 2 or 4GB of memory, a RAID card and a Fibre switch are further options. To accommodate scalable expansion, a variety of rack enclosures, storage towers and PCI-X Expansion Slot kits are available.



GST's D2D Appliance

3. What is a Disk to Disk to Tape Appliance?

Disk to Disk to Tape (D2D2T) appliances are basically an extension of the D2D appliance with the added ability to offload the backup data from the target disk array to tape. Only when D2D is combined with removable tape backup, does it become a comprehensive strategy for protecting against the loss of the enterprise's data under all circumstances, since the backup data that was moved to disk must be moved onto removable media for quick removal and safe remote storage.

The tape can be a single drive, autoloader or tape library. D2D2T has all the benefits of D2D with the added function of offload to tape. Transfer of backup from disk to physical tape media is scheduled at times selected to avoid impacting application performance, making tape speed less important. D2D2T brings together the best features of tape automation and disk in a single, easy-to-use, highly functional appliance.

The GST VirtualDR™ D2D2T appliance provides D2D2T technology integrated into an appliance style package with software included, delivering affordable and fast file replication and recovery from disk, with long-term retention and secure archiving to tape. Tape technologies supported include LTO3, LTO2, and SAIT1 at transfer speeds up to 576GB an hour and up to 13 terabytes of capacity.

GST's D2D2T appliance integrates disk storage, tape automation, network connectivity, and server-class processing with powerful replication and file management software, encased in a single enclosure.

The GST D2D2T appliance tape storage offers low-cost, secure archiving and removability not available with disk-only backup systems. The single-box architecture also eliminates network impact from the migration of data between disk and tape. Like many of the D2D2T appliances, GST's VirtualDR™ D2D2T appliance can backup a wide variety of servers. They connect natively with a FC (Fibre Channel) interface.

GST's D2D2T Appliance





Logically, D2D2T has two separate parts:

- **D2D** – Backing up data from a server's source disk to the target disk array.
- **D2T** – Backing up the target disk array to a tape device.

This requires two separate software resources:

- **D2D** – The existing backup software can be used with few if any changes.
- **D2T** – D2D2T software offloads data from the target disk array to the tape backup device. GST uses its LPAReplication™ software for that purpose.

4. What are the general benefits?

D2D and D2D2T appliances offer tremendous improvement in the traditional disk-to-tape backup and restore process. Both of these technologies provide, at a minimum, the following benefits:

- **Speed** – Disk data transfer speeds have always been faster than tape. D2D backup speeds run at the rate that the host server can stream the data, resulting in a backup operation that takes much less time and can be accommodated in a much smaller backup window. Older legacy networks with their inherent delays result in data speeds that can cause tape drives to constantly stop and reposition (called back-hitching or shoe-shining) that greatly slows the tape drive performance and increases wear on the drive. Virtual tapes do not have to unwind or unload, and virtual libraries do not have to pick and move tape cartridges. Furthermore, multiple servers can be backed up simultaneously, rather than sequentially as would be the case with tape.
- **Reliability** – While tape continues to make strides in improving its reliability, disk is inherently more reliable, both in the device itself and in the media. Disk drives have a Mean Time Before Failure (MTBF) rating of 600,000 to 1,000,000 hours, while the latest advances in tape are in the 400,000 to 500,000 hour range. The superior reliability of disk over tape is in part due to the fewer movable parts in disk versus tape drive, autoloaders and tape libraries. Moreover, the redundancy of disk and use of different RAID formulas enables continued performance with the loss of one or more disk drives. Besides redundant disk drives, there are redundant power supplies and fans that can be hot-swapped in D2D and D2D2T appliances that further increase their reliability.
- **Scalability** – Data in many organizations is expanding at exponential rates; making it difficult to predict backup capacity requirements for the future. D2D and D2D2T technologies can scale from the Gigabytes range to Terabytes so that wasted investments in equipment are minimized. Instead, larger disks can be added or additional RAID modules.



StorFacts™ Report

D2D & D2D2T

- **Convenience** – Offloading to tape can be done later at a more convenient time.
- **Security** – Data that was initially backed up to the target disk array is migrated using a product like GST's LPAReplication™ to move it to tape on a separate LPAR or server. This D2T work is performed in the background of the server being backed up. Because the device looks like a tape drive, the common virus threats to disk storage are not a threat.
- **HSM** – It provides the ability to retrieve data directly from the target disk array at much faster speeds than tape offers, thus offering a form of Hierarchical Storage Management.
- **Reduced Reliance** – Less reliance on having a large backup window to accommodate long backup and restore runs.
- **Storage Admin** – The productivity of Storage Administrators is significantly improved.
- **Less Risk of Failure** – The improved reliability of disk over tape reduces the risk that a backup or restore operation will fail.
- **Backup Software** – Normally, off-the-shelf backup software packages are supported.
- **Tracking** – All copies of data files created are tracked and managed in the backup/recovery library of backup-related files.
- **Lower Total Cost of Ownership** – Efficient design reduces the real day-to-day costs of backup operations, including lower cost to add on virtual tape drives, selective archiving disk cache to tape for lower media costs, shorter operator time for backup, faster recreation of needed costs required less effort, and extended life of tape devices now that the main backup workload taken off of them



Protect Your Data



Many D2D and D2D2T appliances, and all of GST's such appliances are engineered to provide a broad range of storage capabilities, including the management of the storage infrastructure, Hierarchical Storage Management (HSM), archiving, and tape backup and recovery. Major features to look for should include:

- Easily configurable with a wide variety of component choices
- Modular components
- Scalable as needs arise including scalable rack enclosures
- Wide selection of disk and tape hardware, including IBM TotalStorage hardware if a requirement
- Seamless integration with existing backup software packages and policies
- Rack-mountable components
- Low-cost starter configuration
- Field expandability (usually cannot change from D2D to D2D2T)
- RAID support....usually RAID 0 (mirroring), 5 and 10
- Choice of Fibre and SCSI connectivity
- Hierarchical storage ... access to backup data on target disk array

In addition, D2D2T appliances have these additional features:

- Eliminates workload on backup server to produce an extra backup tape
- Improved tape packing rates (the percent of the cartridge filled with data)
- Library management of cartridges and control of all backup data images
- Other library features like an IO slot, multiple tape technologies, auto-inventory checking, bar codes.

5. What is the Hierarchical Storage Management capability?

The backup data that is written from the source server being backed up to the target disk array on the D2D or D2D2T appliance can be processed using HSM. Terabytes of the backup data are on disk as nearline storage and frequently-accessed files can be identified for rapid access. The frequently-accessed files can remain on the target disk for rapid access at disk-to-memory speeds while less-active files are archived off to tape based on established backup rules and policies.

Instead of trying to locate needed historical data on backup tapes that have most likely left the premises, data is more likely to be stored on the target disk array for easy access. In fact, the data files that are most active can be permanently retained on the target disk array, or they can be archived periodically to off-site tape storage once the demand for them has ceased.


6. What types of disks are used?

The Serial ATA II Working Group and the SCSI Trade Association worked together on the interface specs for Serial ATA (SATA) and for Serial Attached SCSI (SAS) so they are compatible today. This permits a level of hardware and software compatibility between the two disk technologies that was not possible in the past.

Some examples of how SAS and SATA are compatible:

- **Wiring** – both use slender 4-wire cables (like telephones use)
- **Cooling** – both benefit from better airflow and system cooling from the slender cables.
- **Connection** – both can be connected by cable or direct-to-backplane; connectors are identical
- **Cables** – same cables and backplane connectors work for both technologies
- **Topology** – same point-to-point topology eliminates jumper configurations
- **Upper Limits** – both are way below the upper limits of their performance limits

Commonalities between the SAS and SATA disk technologies now used in D2D and D2D2T appliances permit them to be used interchangeably in most situations. This adds flexibility to the configuring of the servers in these appliances and permits greater flexibility in reassigning disk drives between appliances and other servers later.



Inside the Disk

7. Optimizing Performance

These new technologies will not automatically provide faster backup throughput. Why? Because just as with tape devices, the server and network have to be properly tuned to assure the speeds stated by the manufacturer of D2D and D2D2T appliances. Storage Administrators need to optimize each link in the backup chain with these appliances, including:

- **Source Disk Array** – The disks on the source/host server backed up must be able to send data at a rate that is sufficient to drive all the subsequent components (listed below) at their fully rated speeds.
- **Media Server** – The bandwidth of the server in the D2D part must not routinely exceed 80 percent of maximum bandwidth during backup.
- **Network Link** – The bandwidth at the SAN or LAN network link must be capable of handling traffic being fed to it during backup.
- **Target Device** – The disk or tape on the target server in the appliance must be able to keep up with the data rate coming at it.



D2D Appliance ~ 750TB



StorFacts™ Report

D2D & D2D2T

Feedback

We value your feedback on this GST StorFacts™ Report. Please send your comments, suggestions and questions to: research@gstinc.com

Related GST StorFacts™ Reports

Click here for a complete list of GST StorFacts™ Reports on data protection.
<http://www.gstinc.com/white/index.html>

About the authors

This GST StorFacts™ Report was prepared by GST's StorFacts™ & Engineering Group under the leadership of David Breisacher, CEO/Chairman at GST. David is the founder of several successful companies, including GST and BCC Technologies, a manufacturer of eServer disk, tape and memory storage devices. A visionary for the storage industry since the early 90's, David lends his market insight and predictions for the IBM midrange storage marketplace to the StorFacts™ conducted at GST. His experience in sensing shifts in technology and industry directions has made it possible for him to organize and structure successful companies to rapidly meet the evolving needs of storage users.

About GST, Inc.

GST Inc. engineers, manufactures, markets and sells a full line of innovative storage solutions including tape, disk and memory products. GST's storage products meet today's need for high-performance, fail-safe reliability and cost-effective data storage protection solutions for the IBM eServer family. This family includes i5/iSeries, p5/pSeries and xSeries servers. GST's advanced tape backup solutions focus on improved backup/restore processes and enhanced disaster recovery. Products range from single and dual tape subsystems and autoloaders to midrange tape libraries and modular enterprise-wide libraries. Unique modular design enables field upgrades, scalability, investment protection and lower life-cycle costs. Development is guided by expert advisory boards that closely track market needs and ensure use of latest engineering technology in product design. Complete information about the company, its products and support can be found at: <http://www.gstinc.com>.

Trademarks

GST, StorageDR, LPAReplication, InternalDR, EntryDR, SafeDR, AutoDR, GrowthDR, ScalableDR, Commander, BridgeLink, StorFacts, SanMatrix and StorMount are trademarks of GST, Inc. in the United States and other countries. AS/400, iSeries, IBM, UNIX, Linux and Windows are the property of their respective owners.