

LPAReplication™ Software



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



Abstract

The increasing need for performance when safeguarding the enterprise's data has led to the development LPAReplication™ technology. This software from GST allows rapid backup from multiple servers to disk for improved performance. Later, backup data is moved off to tape without affecting production.

GST's LPAReplication™ provides the functionality needed to use disk-to-disk backup strategies with a wide variety of eServer and tape device configurations.

This paper explores the importance of disk-to-disk technology as part of the backup function and how GST's LPAReplication solution supports this technology.

GST, Inc. produced this GST Research Report as part of its commitment to provide information leading to improved safeguarding of mission-critical applications and data within the IT industry.



Table of Contents

Executive Summary	3
1. What was the impetus behind disk-to-disk technology?	4
2. How does disk-to-disk technology work?	5
3. What benefits are inherent in disk-to-disk technology?	5
4. How does Hierarchical Storage Management (HSM) work?	6
5. What problems can continue to occur without disk-to-disk technology?	6
6. What configurations are supported by LPAReplication™?	7
7. What kinds of problems are resolved with LPAReplication™?	8
Feedback	9
Related GST Research Reports	9
About the authorship of this paper	9
About GST, Inc.	9
Trademarks	9



Executive Summary

GST continues to monitor developments in the data protection world, talking to IT managers to learn what they are doing with their backup/recovery processes. Many users are looking for more performance than a straightforward server-to-tape backup process offers. In recent years, installations have been turning to disk-to-disk backup solutions:

Disk-to-disk technology first offloads backup data from the server being backed up to a target server; the transfer of data is done disk-to-disk in a highly efficient manner. Then, at a later time, the backup data is moved off the target server to a tape device that permits offsite storage of the tape media for safekeeping.

Advantages inherent in this approach are the greater reliability of backing disk-to-disk (the D2D portion), the greater speed of D2D processing, the faster access to disk-based backup files on the target server, and the control over when to write to the tape backup device. The full disk-to-disk-to-tape process is called the D2D2T process.

D2D2T technology uses separate servers or partitions to permit the initial backup process to be a faster D2D operation. Linux is well-suited to this process.

As a leading provider of tape backup solutions, GST has developed LPAReplication™ software to enable an installation to implement D2D2T with using a wide variety of IBM hardware resources, including i5/iSeries, p5/pSeries and OpenPower servers ... all members of IBM's eServer family. The wide variety of hardware options means that users are able to acquire the benefits of D2D2T at an investment lever appropriate for their needs.

The GST Research Report that follows focuses on the benefits of D2D2T technology as part of a comprehensive data protection strategy for the enterprise. A variety of problems that are addressed by D2D2T and more specifically by GST's LPAReplication™ are covered.



GST's LPAReplication™

The material that follows is meant to provide information that can be useful to organizations assessing the value of D2D2T technology as a part of their backup infrastructure.

1. What was the impetus behind disk-to-disk technology?

The failure of a backup or restore function can threaten the entire organization for days, or even bring it down permanently as demonstrated by the experience of the World Trade Center bombing. Furthermore, as pressure for more production time on servers crowds the time available for daily backups, the need for a faster way to complete backups has become paramount. This initially led to disk-to-disk (D2D) backups that backed up data to another disk at much faster speeds than to traditional tape devices.

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.

In another survey by one of GST's engineering partners of 222 IT managers:

- 66% said their backups take too long.
- 51% worry or know that their current backup schedule leaves data exposed to loss.
- 58% could tolerate no more than four hours of downtime.
- A third experience significant revenue loss or other adverse condition within an hour of application failure.
- Most felt that the anticipation of recovery times was a major problem.
- The growth of ERP data was straining backup windows.

Clearly, organizations are looking for ways to eliminate the drag on server resources by the backup/restore process; and 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 so that backup data can be easily and economically moved to a safe remote location.

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. Thus, the backup data that was moved to disk must be move onto tape for quick removal for remote



storage. We then have a disk-to-disk-to-tape (D2D2T) process that results in backup data being offloaded to tape and moved off to a safe and remote location. That remote location can be a Disaster Recovery Center or a special vault designed to hold and protect backup media against all external threats.

2. How does disk-to-disk technology work?

D2D2T has two separate parts:

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

This requires two separate software components:

- **D2D** – LPAReplication™ or other backup software to collect backup data from one server and send it to a target server.
- **D2T** – LPAReplication™ or other backup software to offload data from the target server's disk array to a tape backup device (tape drive/drives, autoloader, library).

Configuration requirements are quite simple. The target disk array must either be in a separate target server ... or it can be in a separate logical partition (LPA) within the server being backed up.

Therefore, with D2D2T technology: the specialized software is used to backup data to disk in the target server from one or more servers. The target disk array emulates a tape device so the backup software thinks it is backing up to tape. Then, later, disk-to-tape software like LPAReplication™ is used to move the data to the backup tape media for offsite storage.

3. What benefits are inherent in disk-to-disk technology?

Four unique benefits are inherent in D2D2T technology:

- **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 have an MTBF in the 400,000 to 500,000 hour range.
- **Speed** – D2D speeds are far closer to the speed of the server than D2T speeds, resulting in a backup operation from disk-to-disk that takes much less time and can be accommodated in a much smaller backup window, if any is needed at all.



- **Fast Access** – The backup data on the target disk array is much faster to access than data backed up onto tape. GST's LPARepliation software can access the data backed up to disk quickly, offering nearline storage for data that is accessed frequently. The Hierarchical Storage Management (HSM) aspect of LPARepliation™ facilitates the rapid access of files backed up on the target disk array.
- **Isolation of Tape** – The separate D2T process creates the tape media needed for disaster recovery functions that provides a layer of protection when both disk resources in the source and target servers are lost to a disaster. Thus, data that was initially backed up to the target disk array is migrated using LPARepliation™ to a tape device attached to the target server.
- **Security** – The D2T operation is done in the background on the target server, isolated from the applications running on the application/source servers. This isolates critical data from the primary server and minimizes the transmission of viruses. The D2T operation permits the backup data to be removed from the premises to protect it from physical dangers.
- **Ease of Use** – With LPARepliation™, D2D2T is easy to implement and use. Little change is needed in the backup software being used for the D2D portion. GST's LPARepliation™ is policy-driven with menus and fill-in screens that make configuration of the solution simple and self-documenting.

4. How does Hierarchical Storage Management (HSM) work?

Backup data stored on the target server is available for rapid access to the backup files, should the need arise. Instead of locating the data on backup tapes that have most likely left the premises, data can be retrieved from the disk array on the target server. LPARepliation™ contains Hierarchical Storage Management (HSM) capabilities that permit this to be done. In fact, policies can be established for retaining most frequently accessed files on the target disk array for access at a later time. Most-active files can be permanently retained on the target disk array, or purged annually or on some other timeframe.

5. What problems can continue to occur without disk-to-disk technology?

The malfunction of backup processes can quickly escalate to serious or even catastrophic proportions. Successful backup and restore operations are essential to an organization's survival. Without the significant benefits of D2D2T, a variety of problems are made potentially worse:



- **Human error.** Automating the initial D2D backup gets the data protected before the manual intervention processes by the operator of the tape device can occur.
- **Reliability gap.** Most computer downtime events cannot be prevented, only mitigated against. Without the superior reliability of disk as the initial backup, the backup process will have lower reliability measured in MTBF ratings. Moreover, D2D2T adds a full second tier of backup protection not available in traditional D2T scenarios ... the performance benefits of disk are combined with the security and offsite protection of tape.
- **Delayed access.** When backup data is needed in a traditional D2T-only environment, the tape cartridge(s) must be located, returned to the computing site, loaded, read and restored. This process takes orders of magnitude longer than the almost instant access to disk-based data on the target server.

6. What configurations are supported by LPAReplication™?

Although D2D2T technology can operate on many different platforms, it was tailor-made for IBM's eServer series and Linux. The primary server being backed up can have any of these operating systems: i5/OS, AIX 5L, as well as Linux.

GST's LPAReplication™ software runs on the two most popular versions of Linux: Novell SuSE Linux and Red Hat Linux. The server platforms with Linux that are supported by LPAReplication™ are members of IBM's eServer platform family:

- i5/iSeries
- p5/pSeries
- OpenPower

The D2D backup process is conducted on the user's primary servers under their own operating system environment.

The D2T portion on the target server must run under SuSE or Red Hat Linux. The target server can be stand-alone or be located in a separate partition set up for Linux.

This means, from a systems standpoint, to implement D2D2T with LPAReplication™, the user needs to do the following:

1. Set up a separate Linux partition or acquire a separate Linux server of the three types listed above.
2. Install LPAReplication™ on the new Linux partition or server.



GST Research Paper

LPAReplication™

3. Establish data communications (SCSI, iSCSI, Fibre) from the application/source server to the target server.

7. What kinds of problems are resolved with LPAReplication™?

A basic list of data protection needs that can benefit from D2D2T technology are shown below.

- Difficulty in centrally managing different storage resources, OS, cabinet, vendor, connectivity and interface.
- Need to manage SAN and NAS resources together
- Need heterogeneous storage and data center consolidation
- Difficulty in migrating legacy data to new storage
- Protect against out-of-disk conditions
- Compliance with requirements for protection from data alteration/removal.
- Need for connectivity of disconnected SANs
- Require immediate recovery for non-stop availability
- Require remote backup for disaster recovery protection
- Reduce/remove time spent on backups
- Reduce exposure to data not backed up
- Improve storage performance
- Need nearline access to frequently accessed data.

#



GST Research Paper

LPAReplication™

Feedback

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

Related GST Research Reports

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

About the authorship of this paper

This GST Research Report was prepared by GST's Research & 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 research 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. (<http://www.gstinc.com>) engineers, manufactures, markets and sells a line of innovative storage products to meet the need for high-performance, continuous reliability and cost-effective data storage. These products include tape solutions available today, and will include storage-related services, software and disk subsystems in the future. A comprehensive array of tape solutions range from single and dual tape subsystems, autoloaders, midrange tape libraries, to modular enterprise-wide tape libraries, with focus on improved backup and disaster recovery solutions. Modular design enables field upgrades, scalability, investment protection for existing GST tape solutions, and lower life-cycle costs. GST's product development is guided by several advisory boards to closely track market needs and fully utilize the latest engineering developments in product design. Complete information about products, support and company background can be found at the company's Website.

Trademarks

GST, LPAReplication, InternalDR, EntryDR, SafeDR, AutoDR, GrowthDR, ScalableDR, Commander, BridgeLink, 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.