



“ I have been impressed by Coolspirit’s personable and can do approach to everything they do. They have played an important part in ensuring that the customer comes first, that is clearly reflected in what they do ”

Abdo Gadmour
IBM Systems and Technology Group

Our address

24 The Bridge Business Centre
Beresford Way
Chesterfield
S41 9FG

Get in touch

Call us on: 01246 454222
Email us: web@coolspirit.co.uk
Find us: [View location map](#)
Web: www.coolspirit.co.uk

Office hours

mon - thurs 8:30am - 5:30pm
fri 8:30am - 5pm
sat - sun Closed


“

Boost your storage buying power...
use ours!

”

Buy with confidence from
Coolspirit your authorised
IBM Partner





IBM Enables the Data Center with LTO-5 Products — Increasing Capacity, Throughput and File Management

Analysts: David Reine and Mike Kahn

Management Summary

In the animal kingdom, nothing reproduces with the speed and efficiency of rabbits. In fact, there is no such thing as *only two rabbits*. In terms of math, quite clearly, $1 + 1 \neq 2$. If one is a male (*buck*) and the other is a female (*doe*), watch out. Within thirty days you could have a litter of anywhere from four to twelve *kits*. Furthermore, this process can continue multiple times within a single year. Unfortunately, for population control, within only months of birth, the female kits can reproduce, creating even more of a littering problem. Rabbits, in turn, cannot hold a candle to that ultimate reproduction machine, the *Tribble!* In one of the more memorable episodes from the second season of *Star Trek, The Trouble with Tribbles* (1967), we were introduced to these little furry animals that lulled you into a false sense of security with their relaxing trill. Tribbles would have made an ideal house (or starship) pet except for one minor issue: in the words of Dr. McCoy, “They were born pregnant.” **Ultimately, these tribbles consumed all available food resources on the *Enterprise* and filled up all available space on the starship.** *Sound familiar?* It should, if you manage, or even work in, an *enterprise* data center.

Today’s Tribbles are the images, files, and documents that are proliferating throughout the enterprise, uncontrollably. With server sprawl leading to storage sprawl, the data center is awash in multiple, heterogeneous storage arrays. Many of these arrays consist of Tier-1 devices, such as high-performance Fibre Channel (F.C.) disks containing the mission-critical structures that are the lifeblood of any enterprise. Others consist of high-capacity Tier-2 SATA arrays to preserve much of the business-critical and compliance information in the data center. We are also beginning to see an even more performant tier of storage, Tier-0, made up of the highest performing solid-state disks (SSDs) for the most critical table information. In some data centers, the short-term backup of much of this data has moved from tape to disk, using a disk-to-disk (D2D) strategy to take advantage of a more efficient recovery time. For many enterprises, the long-term preservation of this information, however, along with the archive environment, has remained on tape due to the inherent advantages that tape possesses, including its portability, its capacity, and its longevity, up to thirty years when handled properly. This does not even take into account the significant energy savings achieved by using a green tape architecture. However, the amount of data required to be safely preserved is continually growing. In order to maintain a viable backup window, your tape architecture must continue to improve in capacity and performance, as well as security.

The recent announcement of LTO-5 technology specifications by the LTO Program addresses these issues. However, how will LTO-5 products be deployed? With 60 years of experience in the magnetic tape market space, IBM appears to have the deployment issue down pat. With this announcement, IBM is introducing a variety of delivery vehicles to satisfy the tape needs of your enterprise. To learn more about IBM’s LTO-5 offerings, please read on.

IN THIS ISSUE

➤ The Enterprise Data Center	2
➤ LTO-5 Announcement	2
➤ IBM’s LTO-5 Solution Set.....	3
➤ Something New and Really Exciting.....	5
➤ Conclusion	7

The Enterprise Data Center

Over the past 60 years, the enterprise data center has seen a wide variety of architectures come and go, and in some cases, come back and go away again. The *big glass house* has given way to *distributed computing*, which has given way to *scale-up architectures*, which has given way to *client-server*, which now has given way to *scale-out environments*. However, when all is said and done, magnetic tape not only survives, but it thrives, providing the IT staff with the portability and security that it needs to preserve enterprise data and maintain business continuity.

Over the past few years, data center storage has undergone an era of unprecedented growth. Between mergers and acquisitions, consolidations and virtualizations, industry regulations and government compliance, the enterprise data center is storing more data, and more kinds of data, structured and unstructured, such as audio and video, and more copies of data in order to mitigate risk, than ever before. Storage needs typically are expanding by a magnitude of two every twelve to eighteen months. Floor space, energy, and administrative costs are taking a tremendous bite out of the IT budget, not to mention maintenance and other operational expenses. The acquisition costs for much of this storage may be fairly stable, but these ancillary costs contribute heavily to the total cost of ownership (TCO) of the IT infrastructure.

Primary storage will continue to find a home on a heterogeneous mix of disk devices, consisting of the highest-performing Tier-0 SSDs, high-availability Tier-1 Fibre Channel (F.C.), and high-capacity Tier-2 SATA, as will backup images for data with immediate recovery requirements. Enterprise RPO and RTO policies will dictate which backup information needs to reside on disk. However, best practices for data retention in the data center dictate that long-term storage of email and other compliance documents and archiving environments will continue to reside on tape in order to protect the enterprise and its officers from failure to comply with internal policies and government regulations. In a recent study from Fleishman-Hillard Research, it was determined that 70% of companies use both disk and tape to store information, with 53% of them planning to increase their use of tape. **In the end, it doesn't matter where the information came from, it must be protected.** In fact, some data centers that may have evolved to a D2D environment, now find themselves returning to tape in order to take advantage of its high capacity,

Exhibit 1 – Advantages of Tape in the Data Center

- **Lower TCO** – Tape provides the data center with a lower cost for acquisition and operational costs;
- **Energy Efficiency** – enabling the “green” enterprise to fulfill its corporate obligation;
- **Data Security** – through WORM and encryption;
- **Portability** – to enable off-site data protection;
- **Automation** – for high performance and to help eliminate human error;
- **Data Retention** – with up to 30 years of shelf life for the media; and
- **Scalability** – With a space efficient architecture via high capacity in a small footprint.

portability, low-cost WORM, and encryption technologies. These are all innovations introduced on tape - LTO tape - realizing the economies available in terms of energy and other environmental factors contributing to a lower TCO¹, as a result of a reduction in the number of cartridges required. For a full set of advantages available in a tape architecture, see Exhibit 1, above.

LTO-5 Announcement

As with many other elements of an IT infrastructure, tape has seen a variety of architectures come and go over the years. Currently, the accepted winner in the data center for attachment to both commodity and proprietary server architectures is LTO technology. With a decade of history in data protection, LTO technology has proven itself with multiple generations of increased performance and capacity, and even more significantly, multiple sources of both LTO tape drives and LTO media. With January's announcement² of the release of the LTO-5 specification, the tap for tape has been turned on to provide the data center with even more capacity and performance, as well as additional function-

¹ See the issue of *Clipper Notes* dated October 21, 2008, entitled *Disk and Tape Square Off Again – Tape Remains King of the Hill with LTO-4*, and available at <http://www.clipper.com/research/TCG2008056.pdf>.

² See *The Clipper Group Navigator* dated January 29, 2010, entitled *LTO Program Announces Next Gen Tape – LTO-5 Raises the Bar for Tier-3 Storage*, available at <http://www.clipper.com/research/TCG2010002.pdf>.

ality to enhance the long-term storage and archiving capability of the enterprise data center.

With LTO-5 media, the capacity for a single cartridge has increased to 1.5TBs of uncompressed data. This is a fifteen-fold increase in capacity in only 10 years. The native data transfer rate also has increased, from 15MB/s with *LTO-1* to 140MB/s with LTO-5, almost ten-fold. Over that time span, we have seen the inclusion of an integrated WORM capability beginning with *LTO-3* and embedded encryption with *LTO-4*. With LTO-5, we see the inclusion of media partitioning, as part of the LTO-5 specification to improve data management and access through the enablement of self-describing media containers and structured data on tape. (More on this, ahead.)

The availability of the LTO specification simply provides the tape vendors with a blueprint for the delivery of a commodity architecture to simplify the preservation of data and lower the TCO of the IT infrastructure. How that specification is translated into deliverable products is left up to the individual providers. One of those providers is IBM, a company with about 60 years of experience in delivering multiple, successful tape solutions to the IT industry. Its initial response to the LTO-5 specification is no different. It has delivered products to enable increased scalability for data centers of all sizes.

IBM's LTO-5 Solution Set

With the definition of the LTO-5 specification finalized, IBM has proven, once again, to be first out of the starting gate with a variety of standalone and integrated tape solutions in support of an improved information retention environment. Its product set includes a pair of standalone drives, the *TS2250 Tape Drive Express* half-height (HH) drive, and the *TS2350 Tape Drive Express* full-height (FH) drive, along with support of the LTO-5 drive integrated in the *TS3100*, *TS3200*, and *TS3500 Libraries*. All of these products take full advantage of the capabilities of LTO-5 with a native capacity of 1.5 TBs and a throughput of up to 140MB/s, with a SAS interface at 6Gb/s or a FC interface of 8 Gb/s. All IBM LTO-5 solutions also include WORM and encryption functionality, multi-partition capability, and the new IBM *Long Term File System* initially supported in the TS2250 and TS2350 Tape Drives and is planned to be deployed to other LTO-5 products; the latter two are described in later sections. All IBM LTO-5 drives can read and write to LTO-4 (*Ultrium-4*)

media and read LTO-3 (*Ultrium-3*) media helping to protect investments and ease implementation of LTO-5 tape products. Other improvements include *Giant Magneto Resistive Heads* to improve data reliability and extend the life of the drive and the media.

TS2250 Tape Drive Express

Following in the tradition of the *TS2230* LTO-3 drive and the *TS2240* LTO-4 drive, the TS2250 HH external drive provides 1.5TBs³ of native capacity, 88% more than the *TS2240*, in a 2U format, with a new enclosure design where two units can be mounted side-by-side in a 19" IBM server rack, for small-to-medium workload requirements and those data centers with space constraints. In fact, data centers currently still utilizing LTO-3 media will be able to consolidate nearly four LTO-3 cartridges on a single LTO-5 cartridge.

The TS2250 has an increased native data transfer rate of up to 140MB/s, enabling faster reading and writing to help reduce long waits for backup and recovery procedures. In order to improve performance, IBM has implemented twice as many points of speed matching (as compared to the predecessor LTO-4 drive), up to 14, to help minimize back-hitching, supporting native data rates of 40 MB/s up to 140 MB/s. The TS2250 also comes with a 6Gb/s SAS interface, twice that of the previous model.

In terms of new features, the TS2250 supports IBM's new *Long Term File System* to enable the storage of files to self-describing cartridges. It also supports a *SkipSync* feature, also available in IBM's *TS1130* drives, to enable writing to tape without a back-hitch. The TS2250 also has an Ethernet interface to enable IBM Service to update code via an FTP session and an LED to help the IT staff identify when data is being encrypted. Another new feature is *Data Safe Mode* which enables the data center to write protect standard media, providing a logical WORM capability when WORM cartridges are not available.

In addition to the new and improved features, IBM has maintained the functionality available from the previous generations, including:

- AES-256 encryption in GCM mode;
- WORM capability to address compliance;
- Surface control guiding to facilitate the reading and gentle guidance of tape media;

³ 3.0TB with 2:1 compression.

- High-bandwidth dual-stage actuator to help the read/write head to track lateral excursions of the tape media enhancing read/write operations.

TS2350 Tape Drive Express

The TS2350 Tape Drive Express carries the same capacity and data rate specifications as the TS2250, except it is packaged in a full-height (FH) format for medium to large customers with more robust workloads. There are performance differences between the half-height (HH) and FH models in that the maximum tape speed for the FH drive (locate and rewind speed) is 10m/s compared to 6.4m/s for the HH model. There is also a difference in the reel motor acceleration for the two drives, with a rate of 10m/s² for the FH and 5m/s² for the HH. These factors have a direct bearing on the overall job productivity of the two models, with the full-height model delivering more productivity. As can be expected, the FH model is used in the enterprise-level *TS3500 Tape Library* as well as in the midrange TS3100 and TS3200 Tape Library Express Models.

The TS2350 (and the 3588-F5A full height drive described in the next section) incorporates flangeless rollers to help minimize tape edge damage, and a skew actuator to keep the head perpendicular to the tape, all designed to enhance data integrity and longevity. Graceful dynamic braking is also a feature of the height drives to provide a graceful shutdown of the reel motors in case of a power loss.

IBM System Storage TS3500 Tape Library with LTO-5 Drives

The *IBM System Storage TS3500 Tape Library* comes with support for multiple LTO-5 drives, product model 3588-F5A, increasing total capacity and data throughput, and incorporates the same new tape technology features seen in the standalone models. Because of the increased capacity of 1.5TBs per cartridge, the TS3500 can now hold up to 30PBs of uncompressed storage. This can help the enterprise to improve floor space utilization and help to reduce the TCO of the IT infrastructure. In fact, the data center can now support almost 4PBs of LTO-5 data, using a 2:1 compression, in only ten square feet of floor space by utilizing the IBM TS3500 high density frames, model S54.

With room for up to 192 hot-swappable LTO-5 drives, and 20,000 cartridges, the TS3500 can deliver outstanding scalability with an aggregate native performance of up to 26.9 GB/s, helping to tame even the most unwieldy backup win-

dow. Within a TS3500 Tape Library, the LTO-5 drives support a dual 8Gb/s FC interface for resilient, and faster, data throughput.

As with its standalone cousins, the TS3500 LTO-5 drives support data encryption and WORM for data security and compliance. Each drive within the TS3500 can also read and write LTO-4 media and read LTO-3 media, enabling the data center to intermix cartridges to support existing customer data, protecting a significant investment made in multiple generations of media over the years.

IBM TS3200 Tape Library Express with LTO-5 Drives

The *IBM TS3200 Tape Library Express* now comes with support for up to two full-height LTO-5 drives or four half-height drives, increasing total capacity and data throughput, and incorporating the same new tape technology features seen in the standalone models. Because of the increased capacity of 1.5TBs per cartridge, the TS3200 can now hold up to 72TBs of native capacity (144TBs with a 2:1 compression). Available as a standalone library or as a 4U rack-mount deployment, with a standard bar code reader, the TS3200 can help the enterprise to support cost-effective backup and recovery, as well as archival storage, and help to reduce the TCO of the IT infrastructure.

With room for up to 48 cartridges, the TS3200 can deliver outstanding scalability with an aggregate native performance of up to 560MB/s, helping to shrink the mid-size data center's backup window. The TS3200 Tape Library supports an 8Gb/s FC interface, as well as a 6Gb/s SAS interface, for improved data throughput.

As with the standalone drives, the TS3200 LTO-5 drives support data encryption and WORM for data security and compliance. Each drive within the TS3200 can also read and write LTO-4 media and read LTO-3 media, enabling the data center to intermix cartridges to support existing customer data, protecting a significant investment made in multiple generations of media over the years.

IBM System Storage TS3100 Tape Library Express with LTO-5 Drives

The *IBM TS3100 Tape Library Express* now comes with support for one full-height LTO-5 drive or up to two half-height drives, increasing total capacity and data throughput for the smaller enterprise, and incorporating the same new tape technology features seen in the standalone models. Because of the increased capacity of

1.5TBs per cartridge, the TS3100 can now hold up to 36TBs of native capacity (72TBs with a 2:1 compression). Available as a standalone library or as a 2U rack-mount deployment, with a standard bar code reader, the TS3100 can help the SME to support cost-effective backup and recovery, as well as archival storage, and help to reduce the TCO of the IT infrastructure.

With room for up to 24 cartridges, the TS3100 can deliver outstanding scalability with an aggregate native performance of up to 280MB/s, helping to shrink the SME's backup window. The TS3100 Tape Library supports an 8Gb/s FC interface, as well as a 6Gb/s SAS interface, for improved data throughput.

As with the TS3200, the TS3100 LTO-5 drives support data encryption and WORM for data security and compliance. Each drive within the TS3100 can also read and write LTO-4 media and read LTO-3 media, enabling the data center to intermix cartridges to support existing customer data, protecting a significant investment made in multiple generations of media over the years.

Something New and Really Exciting

*Surprise, Surprise, Surprise*⁴

After more than 50 years of tape history and innovation, you might think that tape is incapable of offering any more big surprises. While recent additions of WORM tape and tape drive encryption were outstanding, and while we expect (and LTO technology has not disappointed) that densities, capacities, and speeds will increase continuously, we did not expect tape to redefine itself yet again in another significant new way.

With LTO-5, tape enters another dimension, quite literally. We all have been taught and conditioned to see tape as this singular piece of serial storage. Metaphorically, we envision it as a single piece of physical tape that has a beginning and an end, with “stuff” stored linearly in between at distances that can be measured. We understand that there may be something at the beginning or end of the tape to give us some metadata about what's on the tape (and at what linear location). That's what we have understood ... until now.

What makes tape different from disk? The classic answer is that tape is a “serial” device and disk is a “random access” device. While one can go to an object stored somewhere in the middle of

a tape, its beginning is measured as a distance offset from the beginning of the tape; from that point on, it's all sequential or serial I/O. Disk, on the other hand, appears to be a totally flexible, being able to deliver files or blocks without regard to their actual placement on the disk's physical platter(s).

For most of us, we live in a world of files (of many types), whether documents, music, or programs. We think that we know where they are stored (like the C: drive of our PC) but we really have no idea where or how the file is physically placed. Much of the time, a file can be stored in pieces (called fragments). When we “defrag a disk”, we put the pieces back together again physically, so that the file can be retrieved serially from its beginning (just like tape). **The bottom line to all of this is that neither disk or tape (or files) tend to be what we physically visualize them to be! All have been virtualized, in various ways, and to varying degrees, to allow us to think of them in a logical representation (that makes sense to humans) rather than a totally physical representation.**

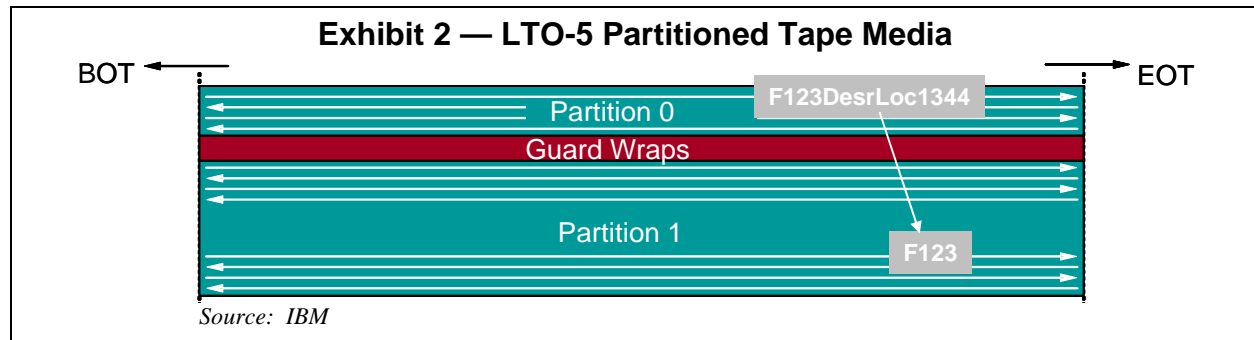
What to Do about Blobs?

Yes, for most of us, our computing experience is all about files. The vast majority of the growth of storage is due to files.⁵ Not only are there now many more of them, many of them are getting quite big. Full length, high-def videos are the most familiar example, where an uncompressed studio version can be measured in tens of gigabytes (or much more, for rendered objects, such as animation). However, there are other kinds of data that also are quite large per object, such as input from a radio telescope, seismic data, and medical images (like MRIs). For decades, computer scientists have called these large, binary (think, non-textual) objects “blobs”. **For many enterprises, what to do about blobs is a big challenge, especially when you may want to keep them for a long-time, often forever.**⁶ Blobs usually are stored as files, often on some form of network-attached storage (NAS). The problems with traditional (disk-based) NAS are:

⁵ Operating systems (like *Windows* and *Linux*) and many data-intensive applications (like some database management systems) manage storage more directly and closer to the physical reality. These control “blocks” of storage and manage what is stored in them (which could be files) and how the physical-to-logical representations (mappings) are handled.

⁶ There are two reasons that something is kept for decades or longer. It may have cost a lot to create (animation) or generate (seismic data) and it may be valuable (and possibly irreplaceable) sometime in the future (like stellar images). Or, it may be legally required, as with retention of medical records.

⁴ Thank you, Gomer Pyle, for this near infamous exclamation.



(1) the arrays consume data center space and energy, wasteful especially for data being held indefinitely; (2) the arrays tend to be replaced every three-to-five years, creating transitional and other challenges; and (3) you need a lot of them, if you have a large collection of blobs.

OK, you might see where this is going. **Storing and accessing files on tape is the next big thing and it could spawn many killer applications.**

LTO-5 Dual Partitioning

Traditional tape cartridges may have an index at the beginning or end. The ones at the front tend to contain well-collected data that doesn't change, because you would have to rewrite the index each time something on it changed. Thus, most indices are at the end of the tape. When new information is added to a tape, it is written at the end, followed by a new (updated) index.

However, the index at the end means that you have to traverse everything along the way just to read in the index, which significantly adds to the access time. It often is as quick to inventory all that you itemized along the way. What you need is a better indexing mechanism.

Consider your nearby Interstate highway. If it's like the ones near us, there are many new markers spaced two-tenths of a mile apart. This is a linear, positional reference, measured, probably, from the beginning of that highway in your state. While it is nice to know where you are with respect to the border, it would be nicer to know what is (nearby) on this highway. While a fixed sign might say, "Dining, one mile ahead", that is only a little bit informative. A highway sign showing what restaurants are at the exit ahead is better but how often is this information updated? An electronic billboard, however, (or your in-vehicle GPS system) likely would be current and might help you find what you seek. These serve as an up-to-date index on what's on (adjacent to) the linear highway. You could search the exits to see what is nearby (i.e., you

could build your own index) but that would be a nuisance. You really want to have the current, usable index available to you, to facilitate finding what you want.

Usable access is now possible with a new feature on LTO-5 media, which can be deployed with two partitions, and using new software from IBM, that turns one of the partitions into an index, with the other partition being used for the data, in particular, large files. The partitions can be variable in size. For instance, one partition can be thin (holding index information) and one can be fat (holding content data). For those who want the details, the thin partition is "Partition 0" and the fat one is "Partition 1". (See Exhibit 2, above, for a pictorial view.) With LTO-5 media, you have a choice: a standard cartridge can be utilized with a single partition, as has always been the case, or "dual-partitioned" (i.e., with two partitions).

What do you do with this partitioned tape? The LTO-5 specification just lays out the partitioned highway. As noted, the size of each partition is variable. It's up to the storage vendors and application developers to decide how to use the partitions. **So, while the dual partition is cool, technically, it is IBM's software that makes this so interesting.**

The IBM Long Term File System

IBM has chosen to use the partitioning as part of its newly announced *Long-Term File System*, initially targeted at organizations with many large blobs that need to be retained long term, especially for the media and entertainment sectors. Early adopters will probably be private industries and other segments maintaining large video files, such as those generated by media and entertainment, security, and surveillance cameras.⁷

So, exactly what is the Long Term File System? Basically, it's a self-contained data

⁷ We can envision many other applications for the Long Term File System. This is just the beginning!

format of data on tape and a file system that is an extension to the operating system. It's the ability to write files directly to tape without a separate, special purpose application to move data to and from the tape. It can put the stuff being stored in the fat partition and put the meta-data and indexing information in the thin partition. It does this in an application-independent fashion.

Each tape is "self-describing".⁸ All the information needed to determine what is on a tape is contained within the tape. It can be read independently, without a separately maintained single-namespace index and manager, or stored as a database, so that a tape delivered to someone 10 years from now will be able to be read without the need for referential integrity (to what is on other Long Term File System tapes). Therefore, it also maintains referential integrity even if that database becomes corrupted; just rebuild the index from the tapes in the library.

IBM's Long Term File System looks and feels like a file system – in fact, it is a file system. Basically, you can mount a tape as if it was a hard drive. It's just that the media is different and the time to retrieval is slower than when retrieving from disk. And, of course, it costs much less than storing the many files on disk and consumes significantly less energy. **This tape-based file system can keep on saving money each year. It enables more to be stored for less.**

For video applications, for example, Long Term File System can offer these potential benefits.

- Reduced cost of video storage media and equipment
- Support for file-based workflow, increasing data mobility and unifying across vendors to a single storage media for video formats
- Reduced data ingestion time
- Reduced archive storage space
- Support for green initiatives by reducing energy consumption

The IBM Long Term File System, enabled by LTO-5 tape technology, is the beginning of something *really* big. It's time to think out of the box!


Conclusion

After all of these years, after all of these decades, tape continues to offer an optimized

balance of both TCO and performance. With LTO-5 tape technology, tape provides the ideal medium for consolidation and scalability, enabling the IT staff to reduce floor space by reducing cartridge count, while at the same time increasing capacity and reducing the amount of energy being consumed for long-term storage and archiving. **LTO-5 protects the investment that the enterprise has made in its IT infrastructure through a commodity interface recognized in every data center around the world and protects that data through encryption and WORM technology.**

In turn, IBM has taken this technology, which they have helped to define, and combined it with the newest innovations in drive and library technology to deliver all the advantages of tape to the enterprise data center, as well as SMBs, and remote offices. With the new TS2250 and TS2350, and the integration of the LTO-5 drive in the TS3100, TS3200, and TS3500 Tape Libraries, IBM has provided larger capacity and faster throughput tape drives with a more powerful long-term backup and archiving solution for enterprises operating in a heterogeneous environment. With IBM value-adding innovation, the data center can rely on improved overall performance and extended cartridge life. With a standardized tape infrastructure from IBM, the data center can help to reduce infrastructure complexity and reduce the total cost of the IT environment. And, then there are all of the possibilities for using IBM's Long-Term File System.

IBM has nearly 60 years of tape storage experience in simplifying the backup and recovery process in the data center. They have helped to remove the troubling gremlins from enterprises around the world during that time, and from our perspective, will continue to lead the way in the years to come. IBM and its LTO-5 offerings have made tape most interesting. Take a closer look!



⁸ This enables hierarchical directory structure, file names, file properties, metadata files, fast search indexes, domain-specific information, etc.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

➤ **The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.**

About the Authors

David Reine is a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, he joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

➤ **Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)**

Mike Kahn is Managing Director and a cofounder of The Clipper Group. Mr. Kahn is a veteran of the computer industry, having spent more than four decades working on information technology, spending the last 17 years at Clipper. For the vendor community, Mr. Kahn specializes on strategic marketing issues, especially for new and costly technologies and services, competitive analysis, and sales support. For the end-user community, he focuses on mission-critical information management decisions. Prior positions held by Mr. Kahn include: at International Data Corporation - Director of the Competitive Resource Center, Director of Consulting for the Software Research Group, and Director of the Systems Integration Program; President of Power Factor Corporation, a Boston-based electronics startup; at Honeywell Bull - Director of International Marketing and Support; at Honeywell Information Systems - Director of Marketing and Director of Strategy, Technology and Research; at Arthur D. Little, Inc. - a consultant specializing in database management systems and information resource management; and, at Intel Corporation, Mr. Kahn served in a variety of field and home office marketing management positions. Earlier, he founded and managed PRISM Associates of Ann Arbor, Michigan, a systems consulting firm specializing in data management products and applications. Mr. Kahn also managed a relational DBMS development group at The University of Michigan, where he earned B.S.E. and M.S.E. degrees in industrial engineering.

➤ **Reach Mike Kahn via e-mail at Mike.Kahn@clipper.com or via phone at (781) 235-0085 Ext. 121. (Please dial “121” when you hear the automated attendant.)**

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.