

Spectra Logic Raises the Ceiling on Archiving — Lowering TCO for Their Tape Libraries

Analyst: David Reine

Management Summary

How many times has a product failed 24 hours after its warranty was up? How often have you outgrown a product with no means to upgrade it? Some companies even have built in “planned obsolescence” so that the consumer must return to the store, or website, for the prerequisite “fork-lift” upgrade. Some companies, such as automobile and appliance manufacturers, may offer “extended warranties” to protect the consumer from possible long-term failures; however, these warranties come with a price, one that must be paid in advance. It is rare when you find a company that makes a product that can grow with you and continues to be supported. Where is your *Atari* game console today? Electronic game machines come and go; however, today’s newest games do not work in yesterday’s controllers. Then there is *Lego*, a scalable toy for your kids that grows with them and is completely backward compatible. If you need a birthday present, new kits are always available, mirroring the latest entertainment craze from *Cars2* to *Harry Potter*. Spare bricks can always be purchased to enable the architecturally-gifted child to scale his or her model as far as their imagination will take them.

Planned obsolescence and upgradability are nothing new to the CIO or data center manager of any enterprise. They have been battling these issues ever since the first computer was invented. How do you support more workers, how do you execute more transactions with the existing IT infrastructure? How can your data center grow like *Legos*? In many cases the answer is: *it can't!* Newer technology is required. Intel, AMD, IBM, Oracle (née Sun/StorageTek), they have all been continually upgrading their microprocessor server technology. Some, such as Intel, seemingly announce a new technology every three months to support more cores, more users, and more transactions. When do you put a stake in the ground and deploy? What do you do with yesterday’s technology? This same process is also true for storage. With primary storage requirements growing (or exploding) on a yearly, monthly, weekly, even daily basis, older arrays become obsolete quickly because of the lack of scalability or the total cost of ownership (TCO) to maintain them when they are out of warranty. In order to prolong their useful life, it is critical that the IT staff identify rarely used data and migrate it off of primary (expensive) storage and onto less expensive, secondary media. One such media is tape, a technology that has been around for over 60 years and can enable the IT staff to access unlimited storage in the data center. One company leading the way with an innovative technology that delivers unlimited scalability, while managing TCO to the lowest possible levels, is Spectra Logic.

With its T-Finity tape library and *BlueScale* software, Spectra Logic enables the data center to migrate rarely-used data off of primary storage while maintaining near-line access to it. To learn more about Spectra Logic’s tape library solutions, please read on.

IN THIS ISSUE

| | |
|-----------------------------------|---|
| ➤ Preserving Data Forever | 2 |
| ➤ Who is Spectra Logic?..... | 3 |
| ➤ T-Finity Update | 3 |
| ➤ BlueScale Update | 4 |
| ➤ Spectra Logic Maintenance | 5 |
| ➤ Conclusion..... | 5 |

Preserving Data Forever

Do you remember the “KB”, otherwise known as the “kilobyte”? There was a time when a KB was a big deal. Not now, as we have progressed from the kilobyte to the megabyte (MB), to the gigabyte (GB), and to the terabyte (TB). How big is a TB? Well a KB equals 10^3 bytes while a TB is 10^{12} bytes! That wouldn’t be so bad for the enterprise data center if we now weren’t talking about petabytes (PB) and exabytes (EB). An EB is actually 10^{18} bytes. Now that is what I call “Big Data”! Moreover, many enterprises and government agencies already are storing many PBs of data. With existing data growth patterns of anywhere from 50% to 100% every 18 months, many of these data centers will be approaching an EB of data (or more) soon. Managing *only* hundreds of PBs is a major challenge; managing an exabyte or more is an even greater task. Running an enterprise data center to manage that much data requires *exascale computing*, and is becoming known as an *exabyte data center*.

Before you select a storage medium for all of that data, it is crucial that you understand the total cost of ownership (TCO) issues for storage and the usage patterns for your big data. The IT staff needs to plan ahead to determine if its data center can support this demand, from an acquisition cost point of view, as well as for the costs of energy, floor space, administration, and maintenance. If you need to store an EB of data (or more), you have to calculate the TCO before you make your deployment decision. Having a scalable architecture is vital to the operation of the data center in order to reduce the complexity inherent in managing and protecting an EB of data.

However, the scalability of data preservation comes at a price. Operating in a pure disk-to-disk (D2D) environment may provide the fastest response time, but it may be more than the IT budget can handle at exabyte scale. As we have seen in the past, having a multi-tiered architecture with a disk-to-disk-to tape (D2D2T) environment may be your best solution¹. **Far from dead, tape is the storage media that may well enable the data center to remain viable and within budget.**

From where is all of this data growth coming? Every enterprise is keeping more information on

hand for longer periods of time, as the value of that data continues to grow. Data retention policies are being based upon both data value and legal requirements to meet regulatory compliance (such as to address eDiscovery inquiries), fear of litigation, disaster recovery, and backup as well as reference activity. Some data files, such as medical images and digital videos for the entertainment industry, may never be deleted, as they may be useful later. In addition to the media and entertainment, industries with tremendous growth rates, with a CAGR of up to 45%, include, but are not limited to: video surveillance, healthcare and medical imaging, genome research, cloud computing, high-performance-computing, and Internet storage and the cloud service providers; enterprises in these industries are heavily involved in exascale computing.

Some of this information must be removed from primary, Tier-1 storage, in order to minimize the need for additional, expensive storage. It should be archived for future use. An *archive* provides opportunity for offline storage, removing “Write Once, Read Rarely” (WORR) data from primary storage and shortening the backup window requirement. An *active archive* does the same things, except that it also provides on-line access, via a file system, for this rarely used information². Clearly, an active archive is a plus for any enterprise that needs to migrate any amount of WORR data from primary to secondary storage.

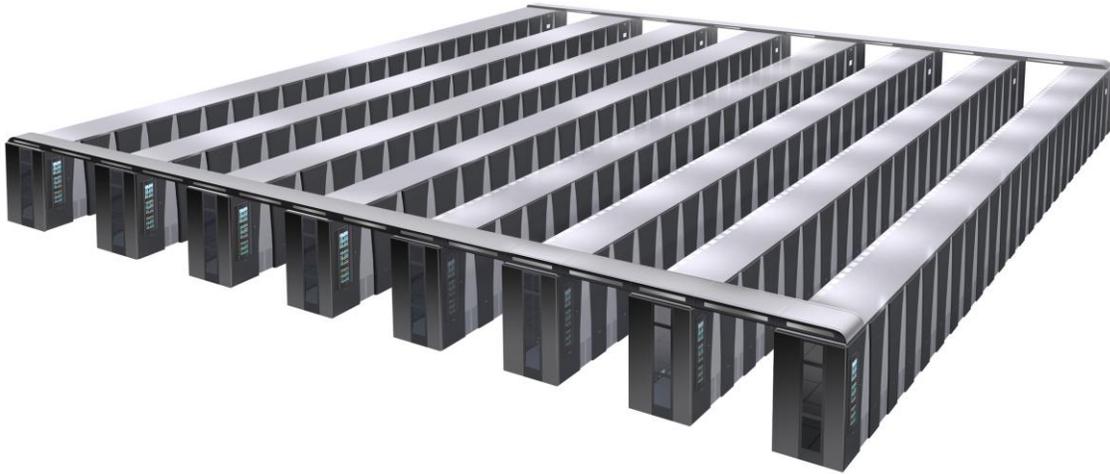
Deployment of a multi-tier environment can improve the performance of your exascale compute requirement by enabling the data center to meet backup window requirements. It will also alleviate cost pressures, as well. The deployment of a tape archive will ease the pressure on the data center budget for primary storage, transferring it to a less expensive medium for data preservation and protection. While “bigger” and “faster” is good for the enterprise data center, the data center still requires improved levels of reliability, density, scalability, and energy efficiency from all tiers.

One company that knows how to do bigger, faster, and more reliable better than anyone, is Spectra Logic. With their recent announcement of added scalability for their T-Finity library, Spectra Logic has raised the bar for archives.

¹ See the issue of *Clipper Notes* dated December 20, 2010, entitled *In Search of the Long-Term Archiving Solution – Tape Delivers Significant TCO Advantages over Disk*, and available at <http://www.clipper.com/research/TCG2010054.pdf>

² See *The Clipper Group Navigator* entitled *Simplified, Online Access to Archived Data – Turning History into an Active Archive*, dated June 11, 2010, at <http://www.clipper.com/research/TCG2010029.pdf>.

Exhibit 1 — Spectra Logic's Skyway



Source: Spectra Logic

Who is Spectra Logic?

Founded in 1979 in Boulder, Colorado, Spectra Logic has been deeply involved in the business of tape storage since 1992, adding tape libraries to their product set in 1995. Throughout the past ten years, Spectra Logic has been a leader in tape library innovation, especially in the areas of encryption, energy conservation³, media lifecycle management, and archiving. Spectra Logic also has been a leader in the formation of the *Active Archive Alliance*, an industry group that was formed to establish archiving standards and develop both hardware and software to facilitate the deployment of scalable archive solutions.

In 2011, in a down economy, Spectra Logic continues to grow in every measurement, adding people, 55,000 sq. ft. of floor space in Boulder and, most importantly, adding to their top line with a growth rate for FY2011 of more than 30% year-over-year as well as their bottom line profits. Spectra Logic also continues to grow in the enterprise space, gaining market share, most notably from Oracle (née Sun/StorageTek).⁴ With 150 customers with at least a PB of tape storage already installed, Spectra Logic continues to make their presence felt in the enterprise data center, as well.

Spectra Logic continues to improve their configurability, scalability, and software functionality. They have improved their tape storage environ-

³ Spectra libraries consume as little as half of the power required by comparable libraries.

⁴ According to Spectra Logic.

ment at multiple levels, from the smallest to the largest, from removing the smallest debris from virgin tapes, to prolonging the life of their tape drives, to deploying the largest tape library available, the Spectra T-Finity, with over 300,000 cartridge slots⁵ and 3.66EB of data capacity.

T-Finity Update

Spectra Logic has made a series of new announcements recently. The most significant, however, may be the expanded capacity of the *T-Finity Library*, now up to 3.66EBs using the *TS1140*⁶ enterprise tape drive from IBM with 4TB tape media. The maximum scalability of a T-Finity library has grown from 26 frames by four rows to 40 frames in eight rows, with multiple robots for high performance and 99.99% hardware reliability for high availability. (See Exhibit 1, above.) This includes redundant components for the robotics, robotics control module, I/O controller, power transmission networks, hot spare tape drives, and power supplies.

To enable the densest possible tape environment, Spectra configures T-Finity with the *Tera-Pack*, a decade-old Spectra Logic innovation to optimize scalability.⁷ A TeraPack contains ten

⁵ Using TS1140 technology.

⁶ See [The Clipper Group Navigator](http://www.clipper.com/research/TCG2011021.pdf) entitled *IBM's New Enterprise Tape Extends Data Retention Capabilities and Lowers the Cost of Data Protection*, dated June 6, 2011, at <http://www.clipper.com/research/TCG2011021.pdf>.

⁷ A TeraPack supports LTO cartridges for the T-950 and LTO and TS1140 cartridges for T-Finity.

LTO cartridges per storage shelf⁸, changing the density paradigm for tape library storage.⁹ For T-Finity, this means up to over 300,000 TS1140 cartridges with up to 960 drives, with encryption built-in for added security. This provides a maximum throughput of up to 2.15 PB/hour for a full complex. T-Finity **supports three times the capacity of other manufacturers, density essential in supporting archives with hundreds of PBs. With a TeraPack, T-Finity can move ten cartridges at a time, with a single robotic movement. T-Finity delivers high capacity, high density, high performance, and extreme flexibility.** The IT staff can configure T-Finity in any combination of master, media, and drive frames to fit your enterprise requirements¹⁰, with service frames at either end to facilitate the importing and exporting of media and to support the hot replacement of robotic assemblies.

When an active archive application is in place to facilitate search and query, T-Finity provides access to all of the historical data of any large enterprise, including media & entertainment environments, high-performance computing sites and government agencies. To facilitate the importing and exporting of large quantities of tape media, T-Finity has a *Bulk TeraPack Access Port (TAP)*, which is also available for the Spectra Logic T-950. The Bulk TAP provides a seven times improvement in the import/export process processing 14 TeraPacks tapes at a time. The Bulk TAP is available as part of the Service Frame in T-Finity, with higher slot capacity than is available in the T950. To facilitate the movement of media between rows within a T-Finity complex, Spectra Logic has developed a unique transfer path for TeraPacks between rows. Their *Skyway* provides an innovative overhead method to transfer TeraPacks between library racks. (See Exhibit 2, at the top of the next column.)

Spectra Logic has also developed a new method for extending LTO drive head life and media life with its innovative *CarbideClean* technology, a new feature of Spectra's Certified Media. CarbideClean uses a carbide blade to pre-clean new media and eliminate microscopic debris that is often left as a result of the manufacturing process. The media is CarbideCleaned prior to shipment and is free of charge to all T-Series customers purchasing Spectra Certified media.

⁸ A TeraPack hold nine TS1140 cartridges.

⁹ A TeraPack can hold nine TS1140 cartridges.

¹⁰ T-Finity supports a maximum of four drive expansion frames.

Exhibit 2 — Spectra Logic's Bulk Tap



Source: Spectra Logic

BlueScale Update

With BlueScale 12, Spectra Logic delivers more than just high capacity, high density, and high performance. BlueScale 12 library management software delivers high reliability and ease of operation as well. With BlueScale, the data center has a single interface with a color touch screen and an intuitive GUI to simplify all management functions for an entire library complex, including setup, configuration, upgrades, security, and reporting – both locally and remotely. This includes encryption key management with *BlueScale Encryption*, fully integrated in the library, and the management of all library partitions. With *BlueScale EnergyAudit*, the IT staff can monitor real-time energy consumption, dynamically. With all of the features built-in, T-Finity does not require multiple external servers for library management, thus lowering energy consumption. Spectra Logic can also be proud of their energy consumption per TB, because of the density provided by their TeraPack configuration.

BlueScale 12 comes with built-in lifecycle management for everything from the media to the library and data integrity verification before and after the drive writes to the media.

Spectra Complete Lifecycle Management

- ***Media Lifecycle Management (MLM)*** maintains the integrity and availability of data by monitoring and reporting on more than 30 different metrics throughout the life of each tape. This ensures the reliability of the media and the safety of your data;
- ***Drive Lifecycle Management (DLM)*** tracks expected utilization thresholds and other health variables, proactively notifying the IT staff to replace a failing drive prior to a hard failure.
- ***Library Lifecycle Management*** provides data on the health of vital library components, tracking usage statistics and advising in advance when components may need to be serviced or replaced.

Data Integrity Verification

Data Integrity Verification enables the IT staff to actively check data already written to tape with a *PreScan*, *QuickScan*, and *PostScan*.

- ***PreScan*** ensures that tapes are usable and can accept data. It checks each imported tape to ensure that the tape is OK for writing data.
- ***QuickScan*** confirms that a single track can be read, scanning the length of that track to provide a rapid indicator of data integrity.
- ***PostScan*** checks an entire tape to ensure that all sectors can be read.

Spectra Logic Maintenance

Spectra Logic also provides *SpectraGuard* support for the data center with multiple options for maintenance, some fairly standard, one not so standard. Standard options include a basic next business day replacement, or more advanced options, including next day, same day, and four-hour onsite service. The not so standard option is for *Assisted Self-Maintenance (ASM)*. This enables the enterprise to eliminate delays by stocking replacement parts on-site, in the data center, to facilitate the replacement of select parts by the IT staff. This includes drives, power supplies, robotics, and the I/O module. This could be vital for the data center that requires minimal downtime or where normal support services are not available.

Having an effective maintenance program in place is vital for business continuity. Being able to afford it is another matter entirely. Typically, once the warranty has expired (typically, after three years, price increases will commence. Interestingly enough, the data center can plan on protecting its archive investments for at least ten years, and perhaps longer. Libraries are upgradable with new

drives and new media, extending their capacity and throughput during that period. With many companies, it is impossible to budget that far ahead, not knowing what the rates will be. With a significant investment in the library itself, however, it cannot be replaced on a whim. Spectra Logic has faced up to this challenge with a new program called *PriceLock*, covering all Spectra T-Series libraries. **PriceLock guarantees the cost of maintenance to be held to existing rates, not increasing above the rate of inflation. This enables the enterprise to budget more effectively throughout the life of the library.**

Conclusion

In the age of “Big Data”, “Big Tape” has become a necessity. Tape is not dead! In fact, tape is alive and well and living in Colorado! Ironically, tape, the technology that was supposed to be dead, is probably in the best position to grow because of our desire to create and retain more and more information. While disk, and especially SSD, will grab the glamour roles, archives with tape will hold the single largest chunk of storage in the exabyte data center, or it had better be if you are going to be able to afford to have an exabyte data center.

Is an exabyte data center beyond your horizon today? Will it still be out of view tomorrow or will you be searching for a scalable solution that can achieve that goal?

With all of the capabilities of an active archive, Spectra Logic can provide real-time access to all of the data in the tape library, no matter how old or how big. The T-Finity library scales to meet the archive requirements of the largest enterprise data center, up to 3.66EB, *at present*. With its BlueScale management software, Spectra Logic not only provides complete lifecycle management, but also the ability to verify the continuity of all of the data in the library. In addition, the reliability and serviceability of T-Finity enable the library to remain available on a 7x24x365 basis. Because of Spectra Logic’s innovation, scalability, density, and maintenance pricing, the low TCO of a Spectra library might enable the data center to remain within budget. If you have unlimited archive requirements, you owe it to your bottom line to take a look at T-Finity and the rest of the Spectra Logic product set.



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 Author

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.)***

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.