SPECTRA

DEEP STORAGE EXPERTS

Enterprise Tape Environmentals

May 2015

800-833-1132 • 303-449-6400 • (Fax) 303-939-8844 • 6285 Lookout Road • Boulder, CO 80301 USA

SPECTRALOGIC.COM

Contents

Abstract **3**
Introduction **3**
Environmentals **3**
Tape Drive – TS1150 Technology **4**
Capacity 6
Compression 7
Reformat 9
Power Efficiency 10
Power Sleep Mode 10
TFinity Library **10**
Density Footprint 10
Scalability 12
Partitioning 14
Power Consumption 15
Integrated Security – Encryption Key Management 16
Common Management Interface 17
Common Components 18
Conclusion **19**

ABSTRACT

Tape storage, including Enterprise tape, long thought to be declining, still holds a considerable amount of the world's data. For example, on an annualized basis, LTO tape by itself saw nearly 24 EB of storage media shipped in 2014 with TS11x0 and T10000 series media from IBM and Oracle respectively adding to that total. [1] The quantity of tape storage shipped has grown every year and is a large part of why data center operators have not run out of storage in spite of significant data growth.

While the quantity of storage shipped is increasing, the unit cost of that same storage is declining at an even greater rate. Compounding this price reduction for users is the fact that enterprise tape offers a number of environmental benefits such as power consumption and cost.

INTRODUCTION

Large-scale data growth has been, is, and will be one of the foremost problem generators in the enterprise data center for the foreseeable future. Data growing at rates in the high double digits annually cause space and budget constraints as well as increased administrative time spent solving addressing these issues rather than developing ideas to move business forward and deliver better services to users.

A great to way address many of the space and budget issues is through the use of enterprise tape technology, generally defined as TS1150 or T10000-series tape drive and media technology manufactured by IBM and Oracle respectively. Enterprise tape is far from its reel-to-reel days and matured well beyond its infancy as an automated solution. Today's enterprise tape is fast, deep, reliable, and easy to use. Furthermore, the impressive benefits for your data center environment include vastly improved efficiency and economics; these have become the key advantages of enterprise tape. This paper discusses in greater detail the many ways in which the environmental characteristics of enterprise tape can benefit your organization.

ENVIRONMENTALS

Data center operations that affect the total cost of ownership for a data center are increasingly focusing on environmental factors. Specifications for air flow, power consumption, cooling, and space and management efficiency technologies are becoming nearly as prominent as the specifications for performance and capacity which have historically been the standard. Coupled with an increased emphasis on emissions control and environmental or carbon footprint reduction, data center operators are becoming increasingly aware of, and responsive to, environmental factors inherent in their facilities as well as those provided by the compute and storage products or solutions within their facilities.

[1] Santa Clara Consulting Group. Backup Tape Tracker 4Q CY2014.

Spectra Logic addresses many customer concerns for environmentally friendly, economically efficient, storage solutions by providing an industry-leading, enterprise tape system that helps reduce the carbon footprint of the data center. Enterprise tape systems drive down environmental requirements related to power, cooling, management, and space savings through a number of methods discussed in this paper.

In many cases, speed and performance also require great power. When it comes to storage, that equation generally holds true. Faster spinning disks require increasing amounts of power to drive those speeds and they require additional power to reduce the heat generated by said disks. Fortunately, not all data must be kept on spinning disk. In fact, a great deal of data may be kept on storage mediums that require far less power, yet still provide solid performance, i.e., Enterprise Tape.

Enterprise tape provides a number of environmental advantages not offered by other storage solutions and does so primarily through two major components – enterprise tape drives and automated tape robotics. This white paper provides detailed information about how they deliver environmentally friendly benefits to the end user.

Tape Drive – TS1150 Technology

The world's fastest tape drive is the TS1150 Technology (TS) drive from IBM. Spectra Logic resells this drive for use in its enterprise tape libraries (T380, T950, and TFinity). The TS drive reads and writes uncompressed data at the rate of 360 MB/s or nearly 1.3 TB/hour. When handling compressed data the drive will read and write at 700 MB/s or nearly 2.5 TB/hour.

In contrast, the next two fastest tape drives on the market, the Oracle T10000D and Linear Tape Open Consortium LTO-6, read and write uncompressed data at the rates of 252 MB/s and 160 MB/s respectively. This means the TS drive is 43% faster than the nearest competitor and 125% faster than the current industry standard LTO technology.

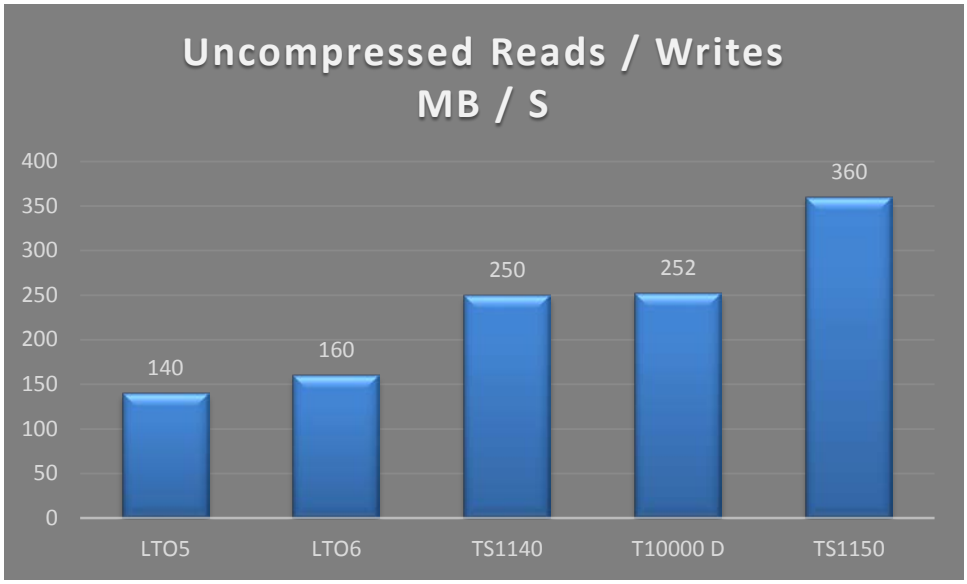

Figure 1: Uncompressed reads and writes for major tape drive technologies

Unlike other storage mediums, the increased performance of the TS drive relative to its competitors does not necessarily come with the power consumption penalty previously noted. In fact, the TS drive consumes far less power than its nearest competitor despite delivering 43% greater performance.

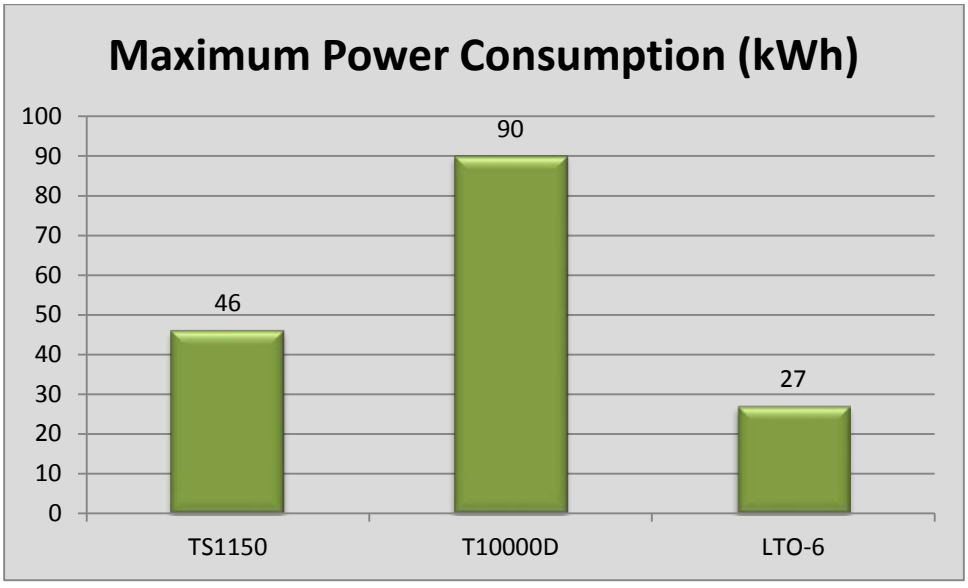

Figure 2: Maximum power consumption for major tape drive technologies

The fact that the TS drives deliver far more performance at comparable or lower power consumption levels results in a tape drive technology that is far more work efficient than the competition. The MB/s read/write performance of TS1150 given kWh power consumption required to generate that level of work performance

makes the TS1150 Enterprise Drive Technology nearly 180% more efficient than the T10000D and nearly 35% more efficient than the industry standard LTO-6 drive.

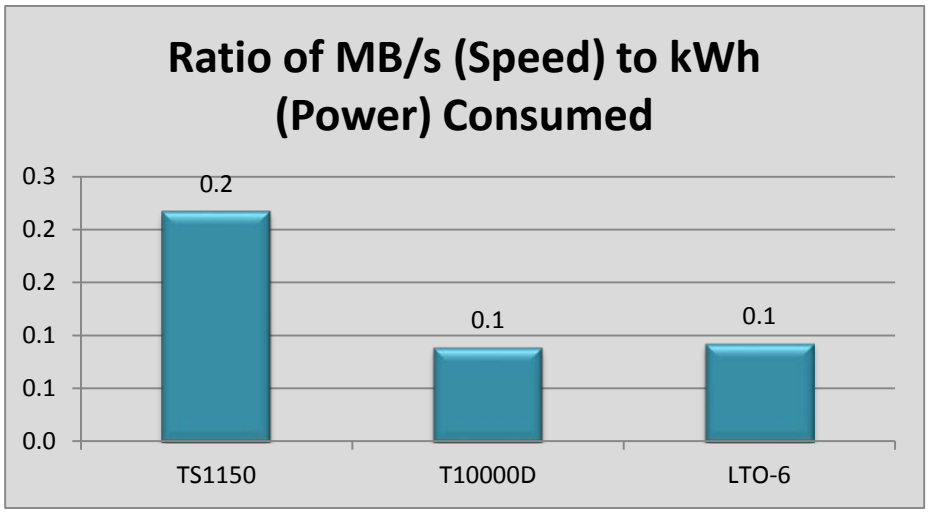

Figure 3: Radio of speed to power consumed for major tape drive technologies

Deploying the most efficient tape drive technology available is only one of many ways in which the TS1150 enterprise tape drive delivers superior environmental performance relative to alternatives.

CAPACITY

Storage capacity can have a significant effect on the environment of a data center. Data continues to grow at torrential rates for many users in which case they must continually add storage. Since storage budgets and data center floor space are finite while the data growth curve seems to be infinite, the capacity and density of storage can either alleviate or exacerbate users' data growth concerns.

Enterprise tape technology like TS1150 delivers one of the highest capacity and data dense storage mediums available today. With a 10 TB uncompressed capacity cartridge, the TS1150 is at the top of its class when it comes to tape cartridge capacity. Compared to the nearest rivals, the T10000D and LTO-6, the TS1150 cartridge provides 25% to 300% greater capacity respectively.

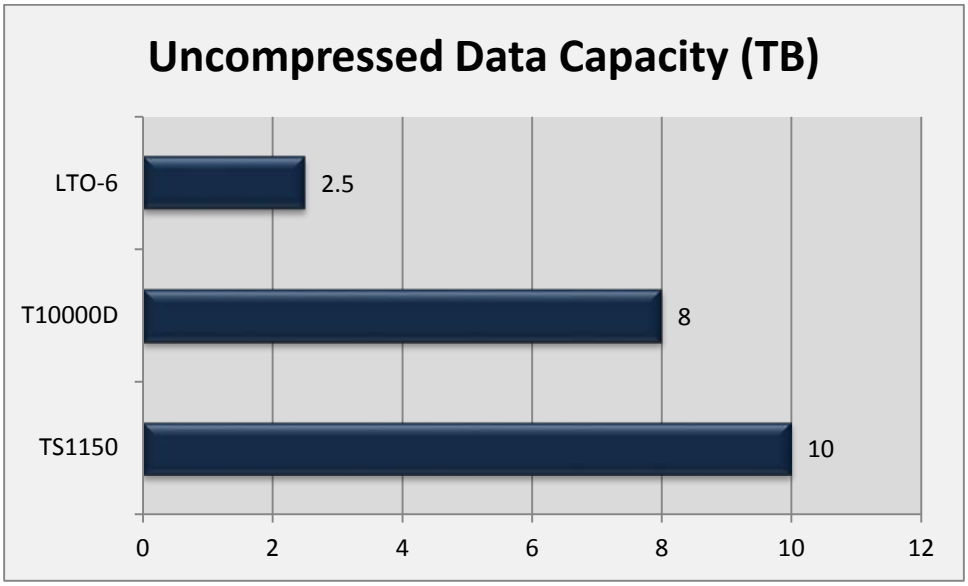

Figure 4: Uncompressed data capacity for major tape drive technologies

COMPRESSION

The data compression feature of TS1150 Technology tape drives further enhances their ability to positively affect the environmental characteristics of an enterprise tape solution.

The TS1150 Tape Drive uses the data compression known as SLDC "Streaming Lossless Data Compression Algorithm – (SLDC)." This method of compression is identical to previous models with the exception of the capacity increase from 1k to 16k in history buffer capacity delivering a better compression ratio than in previous models. SLDC is an implementation of a Lempel-Ziv class 1 (LZ-1) data compression algorithms. SLDC is an extension to the Adaptive Lossless Data Compression (ALDC) algorithm, which is used in leading industry tape products. Users of SLDC can expect to experience the same, or better, data compression compared to users of ALDC.

This means that the TS1150 Technology enterprise tape drive has the ability to compress data at a ratio of 2.5:1. Consequently, the cartridge used with the TS drive, having an uncompressed capacity of 10 TB, is capable of holding 25 TB of compressed data. *For a given quantity of data*, the ability to compress at 2.5:1 means needing over 50% fewer tapes than would be required if compression were not available. Fewer tapes results in fewer tape slots needed and therefore smaller libraries from a capacity perspective. Smaller libraries mean less footprint, or fewer floor tiles of costly data center floor space, will be consumed for storage.

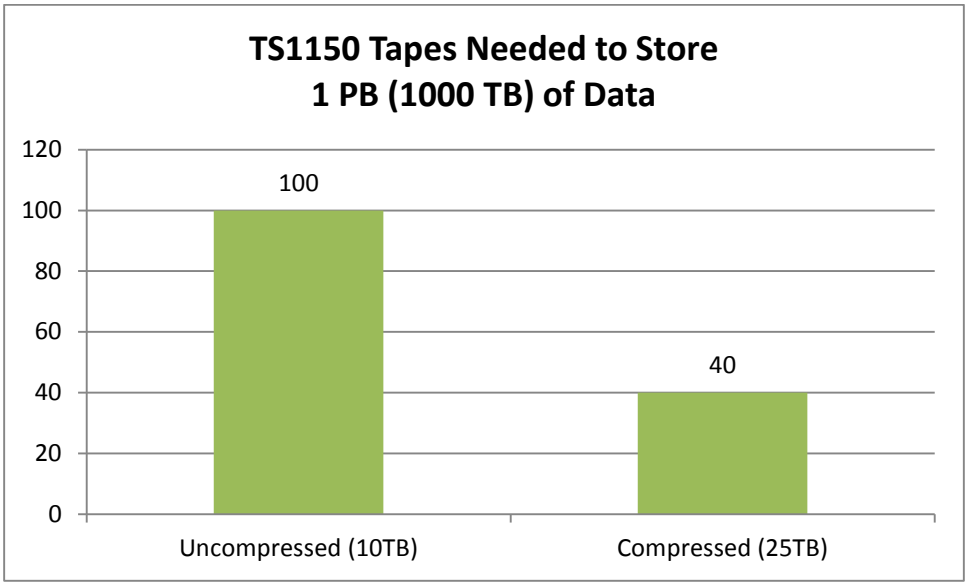

Figure 5: TS1150 tapes needed to store 1 PB of data

Compression delivers environmental benefits beyond just storage capacity. It affects read / write performance in a positive way, as well. The maximum data rate of the TS1150 Technology drive is 360 MB/s, which is the fastest in the industry. With compression, however, that data rate jumps to 700 MB/s. Note that this increase is not quite 2.5:1, due to the limitation of the 8 Gb fibre channel interface. However, the data rate of 700 MB/s is still nearly double the uncompressed rate of 360 MB/s.

Consequently, approximately half as many tape drives operating at their compressed data rate are required to move the same amount of data as are needed when functioning at their uncompressed rate. Reducing the drive count by nearly half, for a given workload, also reduces the size of the library required, helping conserve data center floor space. Additionally, the smaller drive count eliminates a large portion of the solution purchase cost and support contract cost over time.

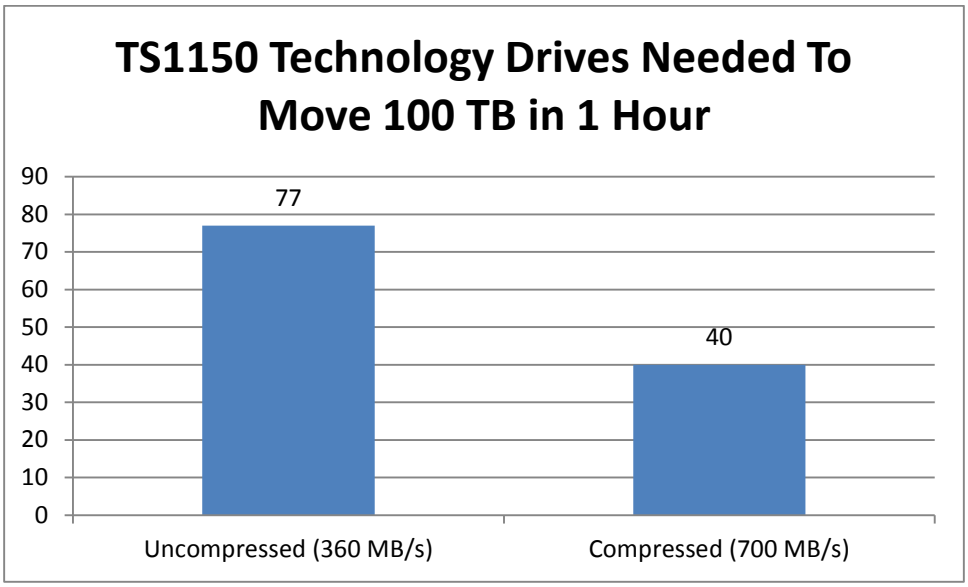

Figure 6: TS1150 drives needed to move 100 TB in 1 hour

The prior two charts clearly show how solid compression ratios on enterprise tape drives like the TS1150 Technology drive will help improve the environmental factors of a given tape automation solution.

REFORMAT

Large capacities and compression are valuable and applicable to current generation tape drives and cartridges, but what about legacy cartridges used with current tape drives? The TS1150 Technology drive excels when it comes to reading or writing prior generation tape cartridges such as those used with TS1140 tape drives.

When a TS1140 tape cartridge is placed into a TS1150 Technology drive, the TS1140 cartridge can be reformatted, starting at Beginning of Tape (BOT), into the TS1150 format. As a result, the TS1140 cartridge will then be capable of containing 7TB of uncompressed data. Since the TS1140 cartridge, written in TS1140 format, is a 4TB capacity cartridge, using uncompressed data, reformatting the cartridge into the TS1150 format increases cartridge capacity by 75% at no additional cost. For users with significant quantities of TS1140 cartridges that can be reformatted, increasing the storage capacity by 75% simply by upgrading the enterprise tape drives delivers a major advantage relative to competitive alternatives and results in similar consolidation benefits previously discussed in the *Compression* section of this paper.

In addition to the added capacity available when reformatting a TS1140 tape cartridge, performance advantages also accrue. In fact, a TS1140 cartridge reformatted by a TS1150 Technology tape drive will yield a performance increase to 300 MB/s which is 20% greater than when used in the original TS1140 format.

Power Efficiency

The tape drive is generally considered the workhorse of an automated tape library system. Its utilization rate can range as high as 80 to 90 percent of the time and possibly even higher in true, high duty-cycle environments. In comparison, the library robotics may be in use only a fraction of the time. As a result, handling its workload reliably and efficiently is the hallmark of a well-made tape drive.

The TS1150 Technology drive stands out in this area. With a power-on consumption rate of 46 watts per hour, it moves the power meter at one-half the rate of the nearest enterprise drive competitor, the Oracle T10000 D tape drive.

Power Sleep Mode

The TS1150 Technology drive provides an environmentally friendly footprint when it's not in use, as well. By transitioning to its power sleep mode when not being called upon, the drive reduces its power consumption from 46 watts to only 22.3. This is a power reduction of more than 50% from normal operating parameters which is less than not only its TS1140 Technology predecessor, but is also nearly one-third less than the nearest enterprise tape drive competitor's standby power mode.

TFinity Library

Tape drives may be the most heavily used component within a tape automation system, but the library itself provides the bulk of the structural environment needed for the total storage solution to function properly. Therefore, it's incumbent upon the library to be environmentally sound if the full solution is to be environmentally efficient. The TFinity Library stands well apart from the competition in many respects, making it the leader in data center efficiency.

Density Footprint

The footprint a library occupies within a data center can have a crucial influence on total data center efficiency. If the footprint is non-standard and doesn't conform to best practice rack row data center layout (as shown in Figure 8), it will adversely affect both air handling and distribution efforts, as well as those of serviceability and optimal floor space utilization.

Spectra Logic libraries, starting with T200 and moving up through the TFinity, are all designed with the TeraPack Architecture which takes full advantage of the height, width, and depth dimensions within the library frame. In contrast, other tape libraries make use of only the height and width dimensions of their interior space. This design difference allows the TFinity, and other Spectra libraries, to enjoy a superior density advantage when it comes to data storage.

Figure 7: TFinity TeraPack Architecture (LTO) and media TeraPack with TS11x0 media

The unique architecture of TFinity and other Spectra libraries use the full height, width, and depth of the library interior, ensuring each library delivers a footprint and density advantage across nearly all configurations relative to the competition. In the example with 5,000 enterprise tape slots and 12 enterprise tape drives shown below, the TFinity demonstrates a significant advantage over the leading competitors in terms of terabytes of data stored per square foot of floor space.

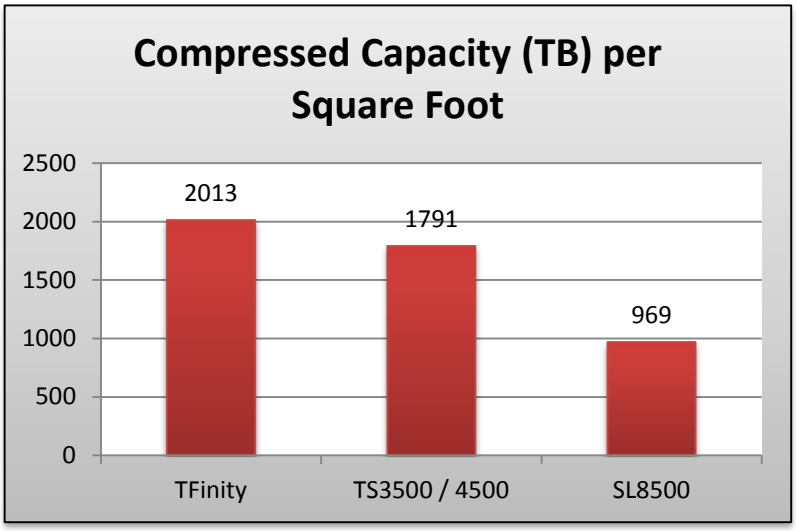

Figure 8: Compressed capacity per square foot

FORM FACTOR

Furthermore, the density advantage of TFinity is being provided within a standard, rack row configuration that conforms to data center best practice. Compared to the Oracle SL8500, TFinity provides not only superior density in a smaller footprint, but

does so without interrupting data center air handling or equipment servicing operations as shown in the comparative example below.

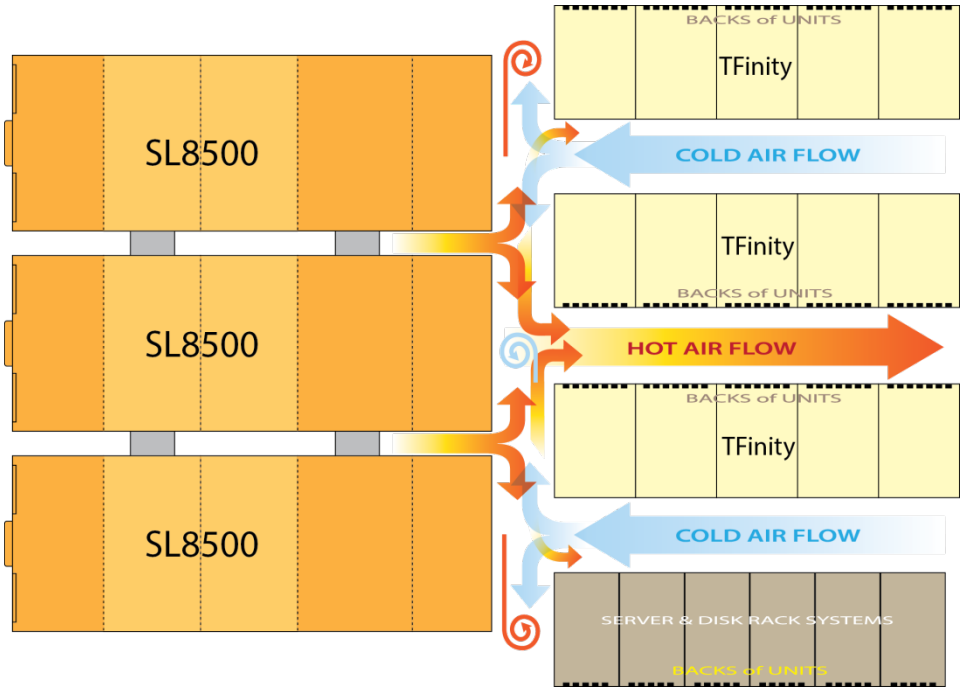

Figure 9: Comparison of SL8500 and TFinity footprint in standard data center layout

The architectural advantages offered by the TeraPack design accrue to both density and standard floor space layout. However, the environmental advantages of this unique Spectra design extend well beyond those of density and footprint.

SCALABILITY

The ability to scale a solution to great size and capacity sounds antithetical to environmental compatibility within a data center. However, considering the economies of scale associated with an ability to grow a solution to large size supports the argument that scalability can and does lead to environmental benefits.

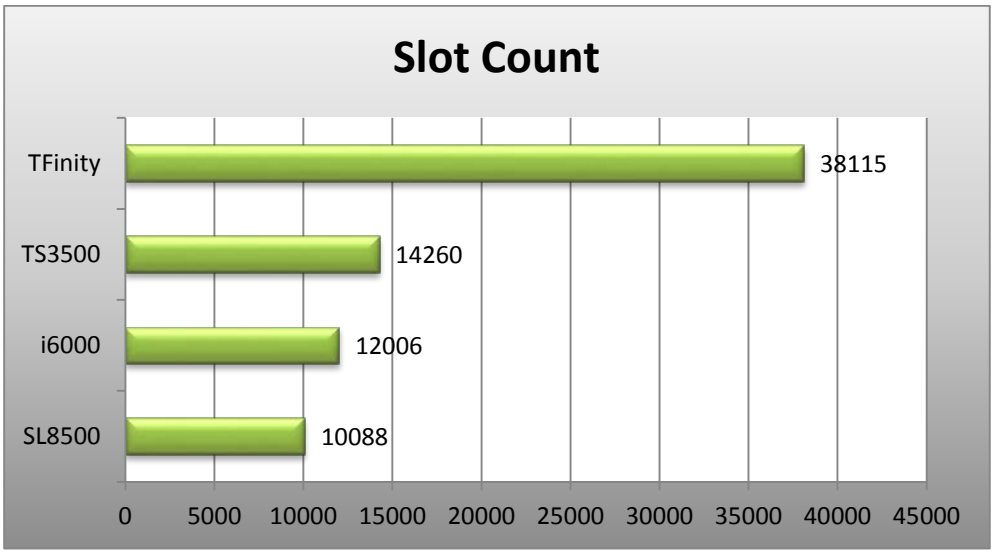

Figure 10: Slot count for competing tape libraries

The graph above indicates the maximum number of slots available per library, using the largest capacity tape technology, with compression, that each library supports. In this case, the TFinity and TS3500 libraries are shown with TS1150 Technology slot counts while the SL8500 is shown using the T10000 D slot count and the i6000 with an LTO-6 slot count. Given these tape technology types, the corresponding storage capacities for the libraries shown are included in the table below.

Table 1: Comparison of slot count and capacity

Library	Slot Count	Drive Technology	Compressed Capacity
TFinity	38,115	TS1150	952 PB
TS3500	14,260	TS1150	356 PB
i6000	12,006	LTO-6	75 PB
SL8500	10,088	T10000 D	161 PB

The TFinity is demonstrably larger than any of its competitors, but that doesn't demonstrate how size equals environmental friendliness. There are a number of ways in which size enables efficiency which include:

- Reduced administrative overhead
- Fewer external servers
- Reduced number of support contracts
- Lower power consumption
- Reduced licensing costs
- Increased storage utilization efficiency

For instance, managing multiple smaller libraries may require multiple administrators. Conversely, managing one larger library can require fewer personnel. The result is that a large library, requiring fewer full time equivalents (FTE) to manage a given amount of storage, allows the extra headcount, otherwise consumed by numerous smaller libraries, to be allocated to other projects thereby improving overall staff productivity. This means getting more work done with a smaller head count.

Many libraries also require external servers to be deployed for things like tape analytics, resource allocation (partitioning), and encryption. The more libraries are required, the more servers may be needed along with their associated externalities. Consolidating as much storage as possible into a single, large library reduces the quantity of external servers deployed. This reduction in server count generally means fewer server support contracts, fewer application licenses, and a reduction in power consumption, all of which improve data center environmental factors.

PARTITIONING

One of the great features of large libraries like TFinity is their ability to be partitioned into many smaller, logical segments or virtual libraries. Rather than purchasing a small library for every department, application, or project it becomes more efficient to purchase a large library and carve it into segments or partitions.

Since each library has individual administration, licensing, security, and resource consumption requirements, multiplying the number of physical libraries in an environment also multiplies costs and complexities. Conversely, purchasing one single library with the capability of being partitioned and shared among users reduces the environmental complexities and costs while still providing users a discrete storage resource.

TFinity can be segmented into as many as 16 discrete partitions with dedicated drives and allocated storage space. The economies of scale relative to deploying 16 physical libraries with individual administration, security, service, power, floor space, and licensing requirements, become readily apparent when a user has to invoke a single set of administrative practices, security policies, storage licenses, and one support contract instead of 16. Economies of scale will also be noted in the consolidation of floor space and reduced power consumption through higher robotic, drive, and storage utilization rates in a shared resource versus lower utilization rates on dedicated resources.

Power Consumption

As the quantity of data grows, so grows the storage to contain it and the power required to keep that storage running. With an estimated 90% of the world's data having been generated within the past 2 years[2], the ability to store the data torrent and power thestorage is of critical importance.

One analyst study concluded the cost of power consumption for disk storage was 238 times greater than that of tape.[3] In fact, the cost of the power consumption alone was greater than the total TCO for a comparable tape solution over a 12 year period in which data growth was assumed to be 45% annually – far short of the rate required to generate 90% of the world's data in the past 2 years. Enterprise tape libraries like TFinity can have a highly positive effect on the carbon footprint of a data center.

Power Distribution

There are several reasons enterprise tape is such a power efficient storage medium. In the case of TFinity specifically, one of the important ways in which it provides such great efficiency is through its internal power distribution system. Robotic libraries have a variety of moving parts, driven by motors, all requiring electricity which must be distributed to those motors through the library.

Many library vendors will install a power supply dedicated to each component that requires power. For example, each tape drive will have its own power supply as will every fan assembly. As a result, the power supplies typically don't have very high utilization rates: they're functioning optimally only when their associated drive, fan, or other powered component is running. Power supplies operating at less than 100% utilization rate are running inefficiently and wasting power.

TFinity's architecture uses a bus distribution system in which a smaller number of larger power supplies are connected to a distribution bus that provides electricity to all the components requiring power within the library. As a result, the shared power supplies are operating at a higher utilization rate and therefore running more efficiently. Because of this greater operational efficiency, and less wasted power, TFinity's overall power consumption is considerably lower than that of its competitors and far below that of disk storage.

An added side benefit to this architecture is realized when a power supply becomes inoperative. The electrical bus system in TFinity, connecting several large power supplies, can provide sufficient power to keep all the components running even when one power supply becomes inoperative. In competitors' libraries, using a dedicated power supply architecture, a power supply failure will generally result in a

---

[2] SINTEF. (2013, May 22). Big Data, for better or worse: 90% of world's data generated over last two years. ScienceDaily.com. Retrieved February 18, 2015 from www.sciencedaily.com/releases/2013/05/130522085217.htm

[3] Clipper Group. (2010, December 23). In Search of the Long-Term Archiving Solution – Tape Delivers Significant TCO Advantage over Disk. Spectralogic.com. Retrieved February 18, 2015 from www.spectralogic.com/index.cfm?fuseaction+home.displayFile&DocID=3712.

failure of its associated drive, fan, or other powered component, thereby degrading the system's overall performance.

ARCHITECTURE

Another important way in which an enterprise library like TFinity delivers great environmental efficiency is through its architecture. Spectra Logic's unique TeraPack architecture, referenced earlier in this paper and used in TFinity, yields much greater storage density per square foot than do either its enterprise tape or disk storage competitors. The combination of high density in a small footprint allows TFinity to improve environmental performance in a few areas.

- Travel distance efficiency
- Work efficiency
- Air handling (cooling) efficiency

The high density, small footprint characteristic of a TeraPack library like TFinity means robots are required to move shorter distances between storage slots and tape drives. The reduced travel distance equates to lower power requirements than would be the case with longer travel distances in larger, less dense libraries.

The TeraPack architecture also enables up to 10 tapes to be transported with a single robot move while 10 individual robot moves would be required in a competitor's configuration. Doing more work in fewer moves improves efficiency and reduces power consumption.

The compact TFinity has a lower air flow requirement for cooling purposes than does its competitors. Requiring a smaller volume of air to be moved a shorter distance to yield an equivalent cooling factor than is required for larger, less dense libraries means the TFinity consumes less power within the data center.

INTEGRATED SECURITY – ENCRYPTION KEY MANAGEMENT

Security isn't normally considered an area in which environmental efficiency is a concern. This is particularly true since data encryption takes place within the storage device, whether it's a tape drive or a disk drive. However, secure management of the encryption keys needed to drive the process at the device level is another matter.

Most encryption key management solutions require one or more external servers dedicated to creating, storing, and managing keys. These servers require additional footprint, power, software licensing, and support contracts. Spectra Logic's libraries, TFinity included, incorporate the key management function within the libraries' internal servers. The key management function is part of the BlueScale library operating system common across all Spectra libraries.

Because the Spectra encryption key management function is internal to the library, the added cost, power consumption, and complexity of connecting and managing an

external server are eliminated. Streamlining the encryption key management process in this way can help improve the environmental characteristics of a storage solution. The simplified key management platform is included within the BlueScale library management interface. The intuitive step-by-step process for generating and applying an encryption key is shown below.

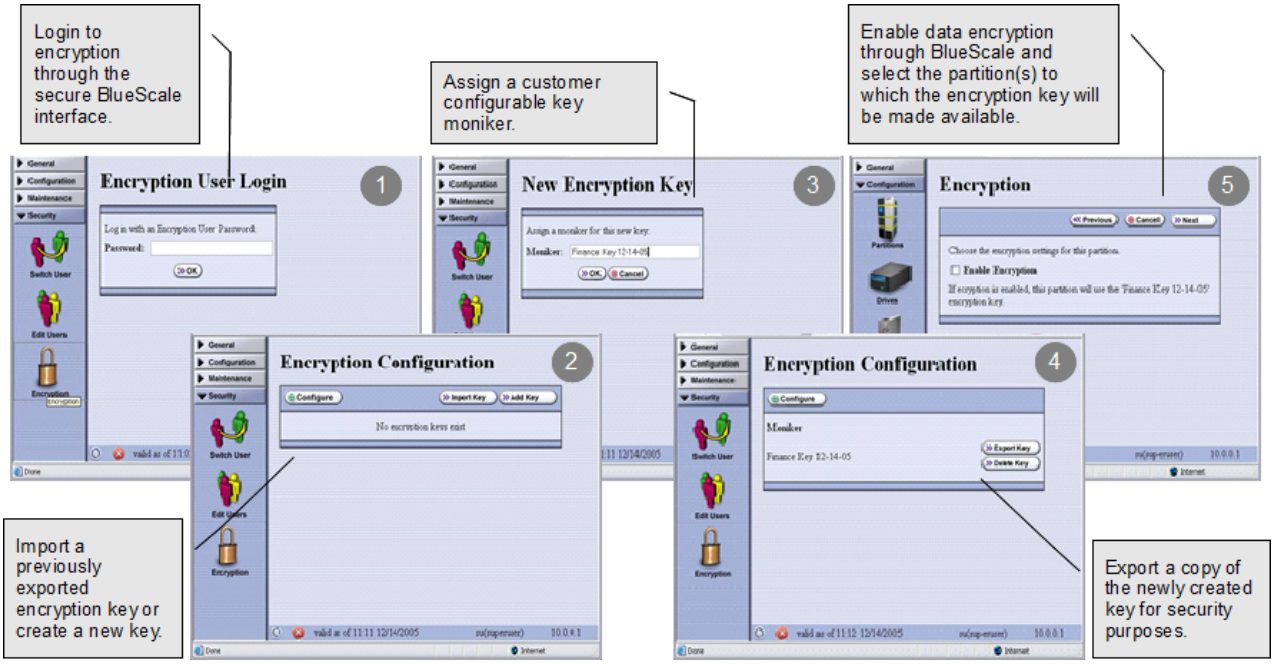

Figure 11: Encryption process in the BlueScale interface

Common Management Interface

Many storage devices, including tape libraries, require multiple management interfaces to conduct all the functions the storage system offers. For instance, there may be one management interface for front panel interactions, a separate interface for remote interactions, and still others for specific functions like configuration, provisioning, reporting, or security.

In some cases, separate management interfaces also require separate licensing, support contracts, operational procedures, and even dedicated servers external to the storage solution. All these requirements add complexity, cost, and potentially additional power consumption above and beyond what's simply required for the tape library and associated drive.

Enterprise libraries that combine these features, functions, and management interfaces into as few touch points as possible help to streamline the automated tape segment of the storage environment. In the case of TFinity, all management functions from configuration to administration to security to reporting have been incorporated within a common library management interface. This is true not only for front panel interactions but also for remote interface purposes. Furthermore,

this common management interface, known as BlueScale, is common across all Spectra Logic libraries.

By combining all these functions within a single, intuitive user interface, TFinity and other enterprise libraries from Spectra eliminate the need for additional licensing, support contracts, procedures, and power consumption due to additional server requirements. The integration of all these functions represents a form of storage consolidation that's unique in the industry, helping drive superior environmental performance.

Common Components

One hallmark of an environmentally sound solution is whether or not it is capable of repurposing its components. Many enterprise tape vendors deliver automated tape solutions that are siloed in that the critical components of one library within a vendor's portfolio cannot be readily used in a different model of library from the same vendor. For customers with various sizes and models of tape automation solutions from one vendor within their environments, this prevents them from taking advantage economies of scale in terms of support options, service practices, and upgradeability.

Spectra Logic recognizes this fact and has moved energetically to address it by architecting its enterprise tape libraries to have a large number of common components. By allowing customers to move, replace, and exchange parts readily among the enterprise library portfolio, users can take full advantage of economies of scale within the support infrastructure of their data centers.

The enterprise family of libraries from Spectra Logic includes:

- T200
- T380
- T680
- T950
- TFinity

These five libraries employ a list of common components that make up the critical operating system within a library. Any of these can be moved, mixed, and matched among the library family. The following list of parts may be readily used within any of the named Spectra solutions:

- Robotic picker
- Robotics Interface Module (RIM)
- Robotics Control Module
- Library Control Module
- Tape drives

- 12v and 24v power supplies
- TeraPacks

Aside from these components, the balance of a library is the external frame and skin which vary with the size of the library. By using common components throughout the library family, Spectra delivers to customers a set of solutions that have the same support inventory, service requirements, and operating procedures. No longer do users need to have a unique set of parts, manuals, tools, diagnostics, and procedures for each library type within their library fleet. As a result, the support and operating structure across multiple libraries is simplified and streamlined.

Furthermore, the common pool of components across libraries enables a fast, easy scalability when moving from a smaller library within the family to a larger version to meet increasing storage demands. In essence, the common architecture helps users streamline their current support efforts and operational activities as well as provide them an easy path to larger configurations in the future.

Conclusion

Enterprise tape has made a lot of progress since its inception many decades ago. It's faster, more economical, reliable, increasingly dense, and readily usable. Furthermore, its effects on your data center in terms of energy, space, and manpower utilization have improved several fold at the same time.

Data growth is transforming the data center environment and generating problems in managing storage space, cost, complexity, and effort. For these reasons, the environmental benefits of enterprise tape cannot be ignored. Rather, they should be boldly embraced as a critical tool in the effort to tame exponential data.

Deep Storage Experts

Spectra Logic develops deep storage solutions that solve the problem of long term storage for business and technology professionals dealing with exponential data growth.

Dedicated solely to storage innovation for more than 35 years, Spectra Logic's uncompromising product and customer focus is proven by the largest information users in multiple vertical markets globally.

Spectra enables affordable, multi-decade data storage and access by creating new methods of managing information in all forms of deep storage—including archive, backup, cold storage, cloud, and private cloud.

For more information, please visit http://www.spectralogic.com.

SPECTRA

DEEP STORAGE EXPERTS

303-449-6400 • 800-833-1132 • (Fax) 303-939-8844 • 6285 Lookout Road • Boulder, CO 80301 USA • spectralogic.com