

Data Reliability: Peace of Mind

The TeraStack® Solution

Executive Summary

New long term data compliance retention requirements, growing data files sizes, longer object lifetimes, and higher resolution video surveillance data have forever changed the scope of data storage requirements. Demanding applications, such as those which store video footage over long periods of time for business intelligence and post event analysis, medical imaging, and data archival require longer data retention and accommodations for rapidly growing volumes of data. With ever increasing importance placed on these requirements, data managers should no longer be content to store longer term secondary or non-transactional data on media with an average lifetime of only 3-5 years. Much of today's large data objects are being stored, not just for months or years, but for many decades.

The TeraStack® Solution significantly refines and refocuses disaster recovery planning goals, video storage management, and data regulation compliance efforts of organizations. The TeraStack® Solution provides significant advantages in reliability at a reduced total cost to organizations, up to 50% savings at implementation and from 60-90% within the first years of operations. Significant cost savings may also be garnered by leveraging the strength of the TeraStack® Solution in the development of Business Continuity and Disaster Recovery Planning (BCP / DRP). A TeraStack® Solution can offer round-robin support of many current backup/archive jobs in place at the primary location, as well as offer a foundation for future scalability through clustering and off-site replication as organization needs dictate. The TeraStack® Solution offers permanent offsite data storage and data redundancy through the inherent reliability, durability, and longevity of Blu-ray™ optical media. The TeraStack® Solution is the combination of the best processor, storage, server, database, and Java Enterprise software technologies engineered to provide the reliable tiered data solution to meet rising data retention concerns.

Introduction

Current long term data retention processes of many companies are based on the expectations and reliability of tape media infrastructure or on large arrays of spinning hard disk drives. Multiple full-copy backup jobs are run frequently to ensure a copy is present on tape in the event of a full disaster. If hard disk drives are used as the primary backup tier, data is immediately mirrored and maintained on additional RAID devices to ensure data integrity. If tapes are used as the primary backup tier, one copy is typically kept on hand and another offsite to minimize the amount of time necessary to locate and retrieve a backup tape for small-scale restore requirements; no true offsite and long term data integrity can be assured due to the expected failure rate and shelf life of tape media and hard drive storage elements. The frequency of full backup jobs can place a great deal of strain on any network and server infrastructure. The large volume of tape media or extra redundant spinning hard drives required for backup purposes places an unnecessary management and monetary burden on the IT department as well.

Hard Drive Solutions

Hard drive technology has advanced rapidly in the past several years, providing higher performance and capacity at significantly reduced prices from years ago, but at what cost? In order to provide the required levels of reliability, hard disks are paired into a Redundant Array of Independent Disks (RAID) configuration for increased performance using striping or for data redundancy using mirroring and parity. Mirroring is accomplished by distributing copies of the data across the RAID to avoid loss of data when one or more drives fail. Striping provides defect management using a single drive or portions of several different drives as error correction buffers. Various RAID configurations can be implemented according to the degree of performance or data protection required. Data managers realize the dire consequences of losing important data with a hard drive failure, and they understand that it is not a question of if the hard drives will fail, but when. In order to protect their data, additional hard drives are being added to the equation, further increasing the workload of data managers and the overall cost of operations. Mirroring the entire RAID with another RAID array offsite for disaster recovery effectively doubles this cost.

The industry norm is to replace complete SAN, NFS, and NAS units before the end of their expected lifetimes, year five or six at best. This is a significant recurring capital and administrative expenditure to maintain the long term archives demanded by today's data requirements, in many cases driven by government mandate. These additional backup requirements, through a mirrored site or via a full tape backup, complicate data management and significantly add cost.

When a drive in a RAID5 or RAID6 device fails it has to be either pulled out, rendering that data unavailable, or it has to be rebuilt while operational, imposing a severe performance cost to the device while the drive is being rebuilt. Should a second drive or even a third drive fail while the first drive is being rebuilt, catastrophic data loss could occur. As the drive capacity continually increases so will the increased probability and amount of data loss within the drives in the event of a significant simultaneous failure of multiple drives. Data managers are fighting to maintain the integrity of their growing volumes of data either by acquiring more RAID devices and commissioning them as dedicated backup devices, or by moving older or less important data to magnetic tape (another mechanical media) for backup or offsite transport for archival.

Hard drives are primarily designed for capacity, performance, and short term availability; not for long term availability and accessibility requirements. For the most part, hard drives are constantly spinning, constantly ready to retrieve the requested data at a moment's notice. The constant spinning, vibrations, harmonics, and heat generated from the myriad of mechanical parts within the hard drives leads to the eventual breakdown of one or more of the components, lending heavily to the 3-5 year usable lifecycle of hard disk drives. Overlooking for the moment the power and cooling requirements and resulting operations costs involved with hard drive solutions, when there is a more reliable, lower total cost storage solution out there; why not use it to the benefit of the organization?

Current trends in society are directing fast immediate access to information. While this may be the case for transactional scientific data, financials, and other transactional or active data, most fixed

data lays unchanged or dormant for many months or even years. In fact, this static or non-transactional sleeping data comprises approximately 80% of all managed data. Long term storage with various hard drive configurations does not fit the mold. Relying on hard drive storage for data rarely accessed is akin to buying a Concorde jet, parking it on the ramp, and leaving its engines running for three years without ever taking flight. Moving non-transactional, aged data automatically onto a high integrity, more energy efficient, and cost effective archival tier represents a lower financial and integrity risk to the organization. Moving static data off of the primary storage tier allows organizations to reclaim extra high speed transactional space for shorter term data, significantly reduce data backup windows, reduce operational costs, and continue the use of existing data center infrastructure without the need for costly expansion.

Magnetic Tape Solutions

Tape based backup and archive has been the backbone of many hierarchical data management policies over the past several years. Its low cost, energy passive media and high capacities have helped magnetic tape manufactures get a commanding hold on the data backup and archival industries. New requirements and technologies are now challenging tape’s hold on these markets. Compliance regulations are sweeping across corporate America, forcing longer term data retention requirements and holding the organizations accountable. Government and medical communities already have long term historical and HIPAA data requirements, but newer regulations, such as Sarbanes-Oxley, OSHA, SEC 17d, and the Federal Rules of Civil Procedure -Rule 26, have greatly extended the data retention periods for which organizations must retain their business data.

Recent legal requirements for data retention have resulted in requirements such as these:

Table 1: Compliance Data Retention Requirements

Legal Statute	Data Retention Requirement
Sarbanes-Oxley	5 years after end of audited fiscal period
HIPAA	Life of patient + 2 years, 6 year min
OSHA	End of audit + 30 years
SEC 17a	End of account + 6 years
SEC 17d	30+ years
15 USC Part 2607	30+ years
US Gov’t Empl Pay Records	30+ years

The Federal Rules of Civil Procedure (FRCP) also specify the timely access of electronically stored information in “reasonably accessible” sources. The burden is firmly placed upon the organization’s shoulders to produce data within a timely manner or have proof as to why the discovery process presents an undue financial burden. Offsite or offline tape archives will create a significant time and cost burden on most companies to comply with FRCP, Rule 26. Not until the tape media is located, delivered, mounted, and spooled up can the discovery process begin. Recent legal proceedings have shown the cost of non-compliance or lost archive data to be increasingly higher, sometimes resulting in significant jail terms and monetary fines, further tipping the scales in favor of longer term, higher reliability media.

Magnetic tape solutions also suffer from poor long term availability and accessibility of data and incompatibility from one magnetic tape format to the next. Although tape manufactures market their product with claims of long life time expectations for high end cartridge tape, common expectations are in the 5 to 7 year range. Many users report that in backup operations, tapes that are used in daily rotation are replaced every six months due to reliability issues. The industry has numerous examples of tapes sitting on shelves for years only to find some degree of data corruption upon restoration attempts. Magnetic media is very sensitive to environmental conditions and especially to outside magnetic fields, such as that created by a magnet or an electromagnetic pulse (EMP) device. Magnetic tapes are also prone to jam, bind, and tear resulting in discarding the damaged tape for a new one. Due to these concerns, data managers will generally create redundant backup tapes, effectively creating a backup of the backup. Magnetic tape data is also accessed sequentially, providing poor suitability for the auditing and searching required during extensive e-discovery litigation requests. Each new generation of tape media has proven significantly incompatible with previous versions and tape readers. This creates a significant financial burden and ties up significant data management resources for organizations when their tape libraries are no longer supported or require migration to a newer version. Although tape reliability has increased noticeably with each new generation of tapes, the possibility of data corruption during long term storage is still a real enough problem to strongly consider more reliable, longer-lived solutions.

Regulations such as SEC 17a and 17d, among several others, also explicitly require write-once, unalterable (also referred to as indelible) media as the only acceptable media to ensure the validity and authenticity of the data. Recent developments in tape and hard drive technologies have provided a means of writing an electronic key onto tape and hard drives to prevent the alteration of data on these media, but given their inherently short lifecycles these are not cost effective solutions for long term data compliance data retention and they are software-based solutions, which are more open to tampering.

The TeraStack® Solution

There is a new approach to long term data management, simple in design, but keyed off of major technological shifts in processing - hybrid storage management, Java enterprise software, and Blu-ray optical media. Blu-ray is ideal in the sense that data, once recorded, can be stored efficiently and reliably and then retrieved for up to 100 years without the worry of lost or inaccessible data. By tracing out the technology shift that this media is now unfolding, along with a core integrated system that optimizes the utilization of this technology, there is a proven approach, with an emphasis on long-term reliability, available to meet the growing demands of data storage management.



(2) units shown in rack



(1) unit shown freestanding


Figure 1: Hie Electronics' TBYTe® 8 Series

Hie Electronics' TeraStack® Solution provides a 78+ terabyte (TB) system, the TBYTe® 8 Series, operating on 600 watts (nominal) power in a 15U standalone or 19" rackmount unit. The TBYTe® 8 Series does not require special cooling requirements or a raised floor data center facility to successfully operate. The TeraStack® Solution meets the growing data reliability void with a total cost of ownership (TCO) savings up to 50% at implementation and up to 90% over a 2-3 year period. This solution is orders of magnitude more reliable than competing hard drive or tape solutions. The product provides a high-end application server running on a multi-core processor coupled with 50TB of (extendable) nearline optical storage, a scalable 28TB (RAW) RAID-capable SATA buffer, and Java-based business rule-based data movement for medium to enterprise class customers in a single compact "data center in a box" solution.

TeraStack® Optical Volumes

The Hie Electronics TBYTe® 8 Series, shown above in Figure 1, can hold up to 8 managed TeraStack® cartridges at any given time; each stack is a 6.25TB (w/ 125 50GB Blu-ray discs) structured optical volume. The TeraStack® cartridge is removable, transportable, and re-mountable carrying its detailed file catalog with it. TeraStack® is a patented system for the retrieval and simultaneous recording of data on optical media. This patented scheme utilizes a reconcilable database snapshot, providing instant access to a remounted TeraStack®. Currently each TeraStack® is stocked with 125 50GB Blu-ray media providing 50TB of reliable nearline, short time-to-access, data.

For long term permanent archiving, when most of the TeraStack® cartridges are recorded, they can be removed easily without stopping the system and new unused TeraStacks can be inserted in their place. Recorded TeraStacks remain indexed on the system and can be easily and quickly brought back on-line for user access. The number of TeraStack® cartridges required is a factor of company data management policies concerning continuity planning, data backup/archival schedules, current data requirements, and data growth rates.



TeraStacks in storage

TeraStack® Features

- 50 to 100 years of data integrity
- Eliminates human contact with media
- 125 Blu-ray optical discs per TeraStack®
- 6.25 TB of storage per TeraStack® with 50 GB discs
- Large capacity flash memory stick attached for master information control system
- 120mm standard backwards compatible disc and burners follows industry road map for future upgrades
- Hot-swappable for easy replacement, long-term storage, and quick upgrades
- Energy passive discs decrease cost of offline storage

Figure 2: Hie Electronics' TeraStack®

Reliability

The TeraStack® Solution was specifically engineered to manage mass amounts of data and to allow access to that data many years down the road. Hie Electronics has chosen Blu-ray discs due to their superiority in reliability and capacity, and long term availability over older red-laser optical chemical dye-based technology, and the Blu-ray industry's aggressive development roadmap for both write-once and rewritable Blu-ray media. Blu-ray media is extremely resilient and guards data against harm from temperature fluctuations, moisture and excess humidity, water damage, power interruptions

or spikes, magnetic variations, critical system failures, and aging. Along with the innate reliability characteristics of Blu-ray discs; Hie Electronics demands the best in quality of construction and operation. Each system is subjected to thousands of use cycles prior to being deemed worthy for operation.

Blu-ray Optical Media

Based on the standard 120mm optical format established for compact discs in 1982, Blu-ray discs achieve superior reliability compared to tape and hard drives mostly due to the chemical alloy structure of the recording media. In write-once Blu-ray media, the heat of the laser at recording time creates a phase change in adjacent layers of copper and silicon from a crystalline structure into an amorphous state which quickly cools forming an alloy that is as durable as the entire disc itself. This high durability alloy will not be affected by high temperatures until such time as the temperature reaches the melting point of the disc. The recording layer of each Blu-ray disc provides very high reliability and life in Blu-ray media.

Durable Protective Disc Coating

To further increase reliability, a clear coating is applied to the recording surface designed to preserve data integrity even after a significant rubbing with steel wool. Manufacturers have proven stable data characteristics even after 10,000 rewrites and over 1 million read cycles. This protective layer provides the following superior improvements over older optical formats.

- Anti-static
- Stain resistant
- Scratch resistant

Improved Error Correction

Blu-ray drives are required by specification to implement error correction techniques to ensure the integrity of recorded data. BD-RE media utilizes error correction to actively verify data and relocate problems to spare areas of the disc from the Data Zone(s) of the disc. Data verification can be accomplished using either one or both of two separate verification phases during the authoring process of the Blu-ray disc. A third level of verification can be added for increased data integrity using a 3rd party application if desired.

Assured Availability

Data availability is the paramount concern for storage of any long term critical data. “Will I be able to access my data 10, 20, or 30 or more years down the road?” Blu-ray media has been subjected to advanced manufacturer lifecycle 80°C/85% relative humidity media testing. Manufacturer testing has shown that write-once (WORM) Blu-ray media can maintain the integrity of its data for up to 100 years. Through this same testing, rewritable Blu-ray media has been rated to reliably store its data for up to 50 or more years. Hie Electronics takes the availability testing one step further by conducting in-house

cyclic testing involving thousands of read-write cycles on each new batch of media to ensure our customers are receiving the best solution for their long term data needs. Hie Electronics is dedicated to providing a solution that will allow operations managers to rest easy at night knowing that their data is safe.

Redundant File Catalog System

The file catalog system for the TeraStack® Solution, used to track where each file is located at any given time, is divided up into several key redundant elements. Several key patent pending designs of the TeraStack® Solution provide the simplified file portability for the 50TB of nearline extendable up to petabytes (PB) of offline data storage. A metadata catalog of file contents on each TeraStack® is stored in three separate locations for redundancy. The master copy of all files is managed within the database management system within the TeraStack® Solution's applications server. Another copy is contained in a small solid state drive attached to each individual TeraStack®. This copy holds a metadata listing of the file contents for its associated stack. A final copy on each disc contains the complete metadata for each individual Blu-ray disc, as well as pointer information associating the disc to its particular TeraStack®. Should anything happen to either copy of the metadata, it can quickly be rebuilt.

Minimized Media Handling

Even though the enhanced coating layer on Blu-ray media provides significant protection from accidental handling, the media should still be handled with care. The media will eventually contain data that is critical to the organization's success and it should definitely be treated as such. Media loading and handling is therefore automated with minimal handling. Straight from the manufacturer, media is loaded directly into each TeraStack® optical volume using an automated loading mechanism. Human contact is only required when transferring the disc spindle from its manufacturer's packaging into the automated media loading mechanism. Once inside the TeraStack, media does not touch other media due to the patented TeraStack® design to prevent unnecessary friction wear and sticking of the media. After the TeraStack® Solution is installed at the customer's location; interior dust is kept to a minimum by utilizing high quality particulate air filters situated directly inside the system's exterior fans.

Pre-operational Stress Tests

All elements of the TeraStack® Solution are subjected to a significant burn-in period prior to commissioning to verify the operational mechanics of the entire system. The robotics, used to move the optical discs, experience a minimum of 10,000 operation cycles prior to commissioning and subsequent release to the customer.

Recoverability

The TeraStack® Solution is also designed with automated self-recovery features when attached to a universal power supply (UPS) backup system. A graceful system shutdown is always the standard operating procedure. The system also has three built in power supplies in a 2+1, active and hot spare configuration respectively. Upon regaining power after a power outage, the system will power up and resume normal operations where it left off. The TeraStack® Solution also has an automatic disc recovery

service that will ensure discs are returned to their proper location after a user configurable period of inactivity. The disc recovery service may also be performed manually by an appropriately qualified system administrator. The disc recovery service helps to recover system resources for other pending read/write operations and to gracefully recover following a power outage.

Long Term Compatibility

The Hie Electronics solution is compatible with the ongoing evolution of 120mm optical discs (CD, DVD, and Blu-ray media) and future optical technology. It has a proven 30-year backward compatibility with standard CD and DVD formats and is designed to work with any future 120mm optical technology as well. A single TBYTe® system supports CD, DVD, Blu-ray, and future 120mm optical formats simultaneously in a single system. All of these formats can coexist in a TBYTe, each with their own TeraStack® cartridge, providing incredible flexibility and growth to utilize existing media and to expand capacity and performance with future generations of 120mm optical media. No longer is there a need to throw out an entire hardware system, and spend critical monetary and personnel resources to acquire a new platform, simply because the media is obsolete.

Significantly Expanded Data Migration Windows

The 3-5 year shuffle of hard drives and the painful 5-7 year migration of data from magnetic tape to magnetic tape are now an issue of the past. Due to the exceptional reliability, availability, and durability of Blu-ray media, its industry standard 120mm format, and its backwards compatibility with older generations of 120mm media and projected future compatibility; data managers are now able to reliably store data for longer periods of time without the worry of being able to access the data when it is needed 20 years or more down the road. Coupled with the added protection and portability provided from each TeraStack® structured optical volume and automated business rule-based data movement, Blu-ray media and the TeraStack® Solution offer a significant ROI to any data storage management environment eliminating costly data migration requirements.

Reliability through Scalability

Significant foresight has been engineered into the TeraStack® Solution to provide simple expansion avenues taking advantage of emerging optical developments and allowing the TeraStack® Solution to scale to meet your growing data needs. The TeraStack® Solution provides excellent system reliability, and simplified expandability and recoverability, through the use of hot-swappable commercial off-the-shelf power supplies, hard drives, cooling fans, optical drives, and other server components, and the structured TeraStack® optical volumes at the heart of the system. The integrated high-end dual core, dual Opteron server and its scalable 28TB (RAW) RAID-capable buffer space provides ample application space for additional best-of-breed applications with access to the entire 50TB+ of nearline data stored in the TeraStack® optical volumes, the online buffer, and any other data accessible within the Enterprise network.

Versatility

A TeraStack® cartridge can be removed from a TBYTe® unit and inserted into another TBYTe® or replaced and stored locally or offsite for satisfying any comprehensive DRP solution. The TeraStack® Solution can contain mixtures of write-once stacks and rewritable stacks; allowing a single TeraStack® Solution to fulfill backup and archive solution roles simultaneously in one solution configurable according to business policies. For instance, day-to-day video surveillance footage could be recorded on rewritable media while predetermined alarm and recording events can be permanently archived onto write-once media for further forensic investigation and possible litigation support. When handling organizational data, the TeraStack® Solution can be configured to rotate rewritable media volumes to provide a planned incremental backup schedule while weekly, monthly, and/or quarterly full backups could utilize write-once media for guaranteed data integrity. Multiple TBYTe® systems will also scale seamlessly together acting as an individual cohesive data solution. TeraStack® optical media volumes are contained in casings designed to prevent accidental damage and unnecessary handling due to transportation or during storage.

Accessibility

Administrator controlled multilevel password and user configuration control is built into the TeraStack® Solution providing controlled access to system resources only to those individuals who are properly cleared to use the system. Also, native encryption schemes already present on the data written to Blu-ray optical media will be preserved throughout the writing process. If additional levels of encryption are desired, other 3rd party encryption schemes may be applied to the data when it is being written to the Blu-ray media.

Data Integrity Statement

Due to the innovative data flow architecture of the TeraStack® Solution, data flows from the data generating application into its file system utilizing the online buffer built into the TeraStack® Solution. Immediately upon the data being stored and grouped into an archive set by the Hie Electronics Java middleware, a copy of that data is written to one of the Blu-ray optical discs in the TeraStack®. During the burn or write process, in-line verification is applied. The data is, therefore, carefully verified at a time close to the data being generated. The TeraStack® can be thought of as being a verified backup copy of active application data that is generated in the background while an application is running. In fact, after the TeraStack® Solution has written and verified the integrity of the data, the most recently accessed data will remain in the online buffer and the original file system location for a predetermined period of time according to business rules. The nearline optical and online buffer copies provide two initial copies of the data. This means that the 5% of the data used in the system 95% of the time is available with online/network access response times, while the infrequently accessed data is moved into second tier lower cost storage. The swap in and swap out of data within the TeraStack® Solution is managed using Hie Electronics' patent-pending Java middleware so that older data is rehydrated from the optical media when requested. The ability to generate duplicate TeraStack®

copies, utilizing industry standard backup software, which can then be taken offline (removed), allows for bank vault disaster recovery storage.

Conclusion

Reliable, long term, high volume media is now a requirement for successful litigation compliance, disaster readiness planning, and long term-high volume video surveillance and forensic investigation applications. Managing and storing ever growing volumes of data on magnetic tape and hard drives can be awkward, extremely unreliable, and accessible data on these mediums 10 or 20 years down the road is doubtful at best. Accessibility is the key for storage of any data. Why should thousands to millions of dollars be spent to maintain long term data on unreliable, high-cost, short lifespan media? The TeraStack® Solution simplifies long term, high volume data access and provides the peace of mind that the data will remain intact for many years to come. The TeraStack® Solution is the “Better way to store and manage your data”.

Hie Electronics, Inc.
321 North Central Expressway
McKinney, TX 75070
Phone: 972-542-2327
Email – info@Hie-Electronics.com
Website – www.Hie-Electronics.com