

Balancing the Seven Key Requirements for Today's Enterprise Storage with Zebi Hybrid Arrays

This white paper reviews the 7 essential features of enterprise storage. It looks at hard disk-based architectures and alternatives such as all solid state disk arrays, hybrid arrays in general, and the unique capabilities of Tegile's new generation hybrid array.



Tegile Systems, Inc.
Web: tegile.com
Tel.: (855) 483-4453

1

EXECUTIVE SUMMARY

Organizations and individuals are creating data faster than at any time in human history, increasing the demand for storage media. As we generate more data, we seek to preserve and protect it with backup and replication, driving the demand for storage media even higher. The result is a significant challenge for IT departments, especially those who want to consolidate IT infrastructure with cloud-based applications, virtualization and file sharing.

Server Virtualization, Desktop Virtualization, Shared Databases and File Services not only demand high capacity storage, they also demand high performance storage that can handle high volume transitions and a larger number of concurrent users. Advances in data reliability, protection and efficiency, such as RAID, compression and de-duplication, also demand better performance.

As a result, IT managers need to rebalance the seven key requirements for storage in order to maintain an efficient IT infrastructure that supports critical business operations and applications. These requirements include: capacity, performance, compatibility, usability (fit for purpose), reliability, data protection and value for money.

To date, many of these requirements have been addressed independently. For instance, hard disk manufacturers have done a great job in adding capacity to disks. However, hard disk performance still lags due to the latency resulting from the mechanical disk movements required to access data. The high speed circuitry of Solid State Disks has successfully slashed latency to microseconds and boosted performance, but the cost per gigabyte is considerably higher than hard disk drives.

IT professionals have been hard pressed to investigate, combine and deploy multiple technologies from myriad vendors in an effort to balance their storage requirements.

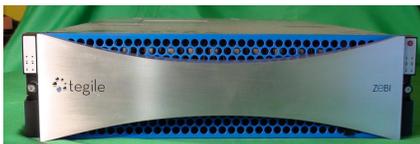
Tegile recognized these challenges and how the seven requirements for storage all need to be balanced in order to support the IT and business needs of the modern data center. As a result, the company is pioneering a new generation of hybrid storage arrays called Zebi.

Tegile's Zebi line of hybrid storage arrays is significantly faster than all hard disk-based arrays and significantly less expensive than all solid-state disk-based arrays. Featuring both NAS and SAN connectivity, they are easy-to-use, fully redundant, and highly scalable. Additionally, they come complete with built-in auto-snapshot, auto-replication, near-instant recovery, on- or offsite fail-over, and virtualization management features.

Uniquely, Tegile's patent-pending MASS technology accelerates the Zebi's performance to solid state speeds without sacrificing the capacity or cost advantage of hard disk storage. Additionally it enables on-the-fly de-duplication and compression of data so each Zebi has a usable capacity far greater than its raw capacity.

The 7 key requirements for enterprise storage:

- Capacity
- Performance
- Compatibility
- Usability
- Reliability
- Data protection
- Value for money



Store more via In-line Compression and Block-level De-duplication

2

STORAGE CHALLENGES

When IT departments are tasked with adding enterprise-class storage, there are seven key requirements that have to be balanced when choosing an array. These are:

1. Capacity – Does the array have the capacity I need and the ability to expand?
2. Performance – Is the array fast enough for my current needs and anticipated growth?
3. Compatibility – Will the array work with my servers, file systems, and network?
4. Usability or fit for purpose – Can the array be configured for the application(s) it's intended to host without adding complexity?
5. Reliability – Can we rely on the array to stay online even if a component fails?
6. Data protection – Does the array have features to help protect my data and comply with government regulations?
7. Value for money – Can we afford the array and if so, are the features commensurate with the price?

Some of these feature considerations are clear cut. The array is either compatible with the desired infrastructure and applications or it is not. Additionally, there's usually a minimum capacity requirement, maximum price restriction, and the expectation that the arrays being considered are reliable.

The other feature requirements tend to represent the competitive value of the array. Over the past decades traditional array vendors have maintained or improved each of these features. Capacity is increased when the storage density of the component hard disk drives is increased. These storage density improvements drive down the cost per terabyte. Performance is improved by arraying more and more hard disk drives together in one system. Short stroking spindles can increase performance somewhat, but it results in wasteful under-utilization of the available disk capacity. Data protection capabilities are improved by adding software.

Unfortunately, when the value of an array is driven by adding hard disk drives and software, it becomes complex, bloated and unnecessarily expensive. That is what has happened over the years. While the architecture of enterprise arrays has changed little, the way IT organizations want to use them has changed radically. IT organizations are increasingly consolidating their compute and storage infrastructure using virtualization technology and cloud-based applications that can be shared by thousands of simultaneous users.

With more users the criticality of performance, capacity, and data protection has increased to the point where the value of purchasing and managing hard disk-based arrays is challenged by new arrays based on Flash solid state disks.



Solid State Disk Alternatives

Flash memory-based solid state disks are significantly faster than hard disk drives delivering the performance required for today's data centers. At the component level, a solid state disk with a standard SAS interface could promise speeds up to 175,000 IOPS – that's 1,000 times the promised IOPS of some 15,000 RPM hard disks that use the same interface.

On the other hand, solid state disks have a considerably higher cost per terabyte compared to hard disk drives. And although the cost of Flash is falling, so is the cost of hard disk drives, thereby maintaining their status for the foreseeable future as the best value for storage capacity. The higher cost per terabyte of solid state disks is exacerbated when organizations also have to pay for the capacity required for data protection features and mirrored systems.

In addition, early adopters of solid state disks have reported significant drop-offs in performance over time. Some systems start out fast, but their performance can quickly plummet to levels akin to higher capacity, less expensive, feature rich hard disk-based systems.

In a nutshell, solid state disks solve performance problems. As such they have their place when only speed is paramount, but capacity, price and data protection features are not.

Hybrid Alternative

An alternative to the slower, bloated, all hard disk-based array or the fast, expensive, all solid state disk array is the hybrid array. The hybrid architecture uses fast solid state disks and/or DRAM to cache in-demand data and uses less expensive hard disk drives for low cost capacity. The result is a faster, high capacity array at a reasonable price.

Hybrids make sense for IT organizations that need arrays that balance performance, capacity and price. However, on its own the hybrid architecture does not necessarily address the "add more to do more" challenge IT managers might expect of a new generation of arrays.

The choice of storage media - hard disk, solid state or a mixture - is important, but not the only criteria. Today's systems must also deliver a host of advanced features.

Advanced Features

The need for data protection to safeguard data assets and comply with regulations is universal. Today's data centers have come to rely on technologies such as snapshots and replication to meet their data protection needs and overcome the limitations of backup windows. Although many traditional and hybrid arrays now offer snapshot and/or replication capabilities, these features tend to carry a performance penalty and often command an additional price premium. Bandwidth consumption can also be a major inhibitor.

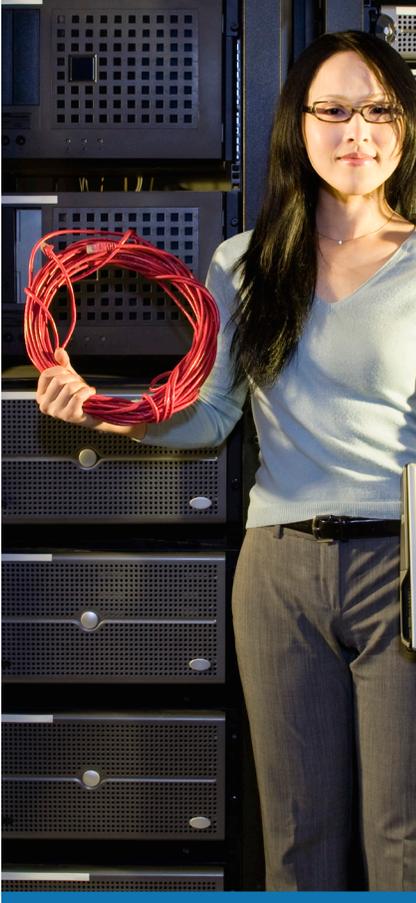
Some arrays also offer compression, which can be very useful in helping to reduce the overall storage capacity requirement. De-duplication is



available from somewhat fewer vendors, again usually at a price premium. Like compression, de-duplication can be a very effective way to increase efficiency and reduce the capacity requirement of storage. However, de-duplication tends to come with a particularly high performance penalty – so much so that it is only offered as a backup or secondary storage option by many vendors since it would be too slow for primary applications.

Today's data centers tend to be diverse, often with a mixture of fibre channel and iSCSI connectivity and a reliance both on structured database-resident data as well as unstructured files. Too often, storage arrays are limited in their application and connectivity. For IT professionals, this leads to complexity, inflexibility and adds significantly to both capital and operating costs as they are forced to manage silos of storage. The increased complexity of provisioning and managing storage in this way can also lead to errors and costly downtime.

Even virtualization brings challenges. Server virtualization puts greater demands on storage to deliver data without latency – clearly a problem for mechanical hard disk drives. Virtual Desktop Infrastructures (VDI), whether based on VMWare View, Microsoft Terminal Services, Citrix Xen or other solutions, notoriously cause disruptive boot storms. Low-latency Solid state disk is only a partial solution as virtualized environments need high capacity, value, and advanced features in addition to performance for their storage.



3

TEGILE ZEBI SOLUTION — A NEW GENERATION OF HYBRID ARRAYS

Most of the mainstream disk arrays available today are based on graying product lines dressed up with incremental modifications. These over-priced systems are based on old architectures, legacy hardware, complex command line management interfaces, and pay-by-feature software.

While mainstream arrays have changed little, the way businesses want to use them has changed radically. Businesses are consolidating their compute and storage infrastructure using virtualization technology and cloud-based applications that can be shared by thousands of simultaneous users. Additionally, IT managers often have special situations where using a mass-produced array is like forcing a square peg into a round hole or they need a little one-on-one support from a storage expert.

Today's IT managers need storage arrays that deliver faster performance, higher capacities and robust data protection with near-instant recovery times. They want these arrays to be affordable and easier to use. Additionally, they want to work with a supplier that listens to their specific needs and provides customer support as an integral part of their offering, not as an afterthought.

Tegile Systems set about becoming the "innovation partner" IT managers really need to deliver a new generation of storage arrays that address all of the demands of the modern data center.

Tegile's Zebi Storage Arrays

Tegile's Zebi line of hybrid storage arrays is up to 5 times faster than all hard disk-based arrays and significantly less expensive than all solid-state disk-based arrays. The usable capacity of a Zebi is up to 5 times its raw capacity because of Tegile's on-the-fly data compression and de-duplication technology. Zebi provides better data protection than standard arrays with its no-single-point-of-failure architecture, built-in auto-snapshot, auto-replication, near-instant recovery, on- or offsite fail-over capabilities. Zebi arrays integrate easily into both SAN and NAS environments. Its comprehensive and intuitive browser-based dashboard, with its configuration wizard and analytical tools, is easier to use and more helpful than other user interfaces.

Up to 5X Faster

Tegile Zebi Array



The Tegile Difference

Tegile's patent-pending technology underpins the company's pioneering delivery of advanced, cost-effective storage for 21st century data centers. Tegile's Zebi array uniquely combines a Redirect on Write (ROW) file system with proprietary Metadata Accelerated Storage System (MASS) technology to realize the holy grail of storage – high performance, high capacity and high reliability at low cost. It combines this unique technology approach with best in class data protection features, the most reliable hardware and an extremely friendly user interface to take storage to a whole new level of usability.

Tegile's highly experienced team underpins its reputation as the IT manager's "innovation partner." When customers need help, real experts are there for them. When a customer provides feedback, Tegile executives evaluate it immediately. When multiple customers provide the same feedback Tegile take immediate action to see how it can benefit all customers. This simple process makes Tegile the ideal data storage partner for business, government and educational organizations.

Metadata Accelerated Storage System (MASS)

Traditional storage systems store data and metadata together, with metadata being interspersed with data on disks. Over time, with data being modified, deleted, and rewritten, metadata becomes very fragmented on disk. Additionally, certain storage system features, such as de-duplication, can cause metadata to multiply and grow rapidly.

Inordinate growth of metadata causes significant deterioration in a system's behavior over time. This was the key driver behind MASS. With MASS, the Zebi storage system organizes and stores metadata, independent of the data, on high-speed triple-mirrored devices with optimized retrieval paths. This accelerates every storage function within the system, raising the performance of near-line SAS hard disk drives to new levels.

"Prime-time" De-duplication and Compression: Tegile's MASS technology not only caches in-demand data to fast solid state disks to deliver high performance, it also enables on-the-fly de-duplication and compression. These data reduction technologies give Zebi a usable capacity up to 5 times its raw capacity while maintaining the array's high performance.

That's a significant development, not just because it solves the "add more to do more" bloat of legacy arrays. The performance of legacy arrays is significantly impacted by de-duplication and compression processes. Tegile's MASS technology mitigates the performance penalty associated with de-duplication and compression allowing data reduction technology to be applied to primary storage, not just near-line or backup storage.

Virtualization Benefits: Tegile's prime time de-duplication and compression capability is a major development that's important to IT organizations implementing desktop virtualization (VDI) and server virtualization. In these environments Tegile's unique capability can result in storage savings of up to 90%. Additionally, the high performance capabilities mitigate the impact of "boot storms" and a growing user base.

*Usable capacity
up to 5X the raw
capacity*

*VDI storage savings
up to 90%*

Snapshot Improvements: Tegile's MASS technology also accelerates the built-in data snapshot feature that allows an IT manager to quickly backup and restore data. Snapshots can be scheduled every few minutes, every hour, day or week depending on the criticality of the data. Snapshots can be restored almost instantly. Additionally, Tegile's unlimited snapshot technology is non-disruptive and extremely space efficient because it records the changed data, without duplicating existing data. And unlike other vendors, Tegile provides its superior snapshot technology as a built-in feature of the Zebi array, not as a costly software add-on.

Other Innovations

The patent-pending MASS technology helps Tegile deliver storage arrays that are up to 5 times faster than all hard disk-based arrays and significantly less expensive than all solid-state disk-based arrays. It also enables Tegile's Zebi arrays to offer a usable capacity up to 5 times their raw capacity. Plus, it helps Tegile deliver superior snapshot and recovery capabilities. However, Tegile arrays include even more innovations.

Zebi arrays provide better data protection than standard arrays with their no-single-point-of-failure architecture, auto-replication, near-instant recovery, and on- or offsite fail-over capabilities.

Tegile also built in true unified access into Zebi. Both block protocols (iSCSI, Fibre Channel) and file protocols (NFS and CIFS) are equally well supported. This makes it possible, for instance, to deploy Zebi for virtualization as a Fibre Channel target to provide high-performance, low latency storage to the hypervisor, while simultaneously using the file share features of Zebi via CIFS for user file shares.

And the Zebi's comprehensive and intuitive browser-based dashboard, with its configuration wizard and analytical tools, make it easier to use and more helpful than other user interfaces.

Unified Storage

Zebi uses a no-single-point-of-failure architecture with hot swap components



4

ZEBI BALANCES THE SEVEN KEY REQUIREMENTS OF STORAGE TODAY

The Tegile Zebi storage array is a multi-protocol SAN & NAS solution that enables consolidation of Tier 1, Tier 2 & Tier 3 storage on a unified storage platform. It effectively addresses the seven key requirements for storage in today's data centers in a uniquely intelligent, balanced manner.

The 7 key requirements for enterprise storage:

- Capacity
- Performance
- Compatibility
- Usability
- Reliability
- Data protection
- Value for money

Capacity

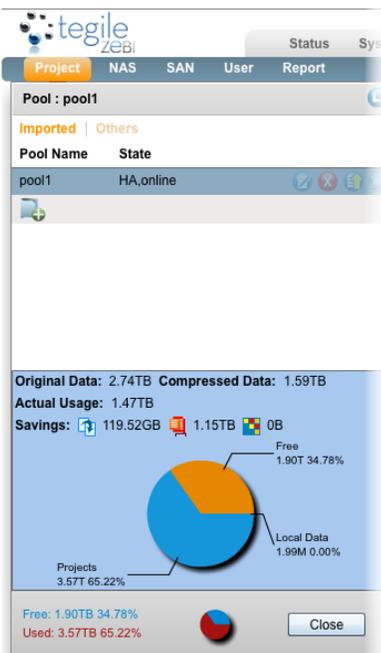
The usable capacity of a Zebi array is up to 5 times its raw capacity allowing organizations to store more data per unit of rack space. Built using enterprise class multi-terabyte hard disks in 2U to 12U configurations Zebi is less than half the size of many storage incumbents, reducing data center space and power consumption. Additionally, thin provisioning and superior snapshot capabilities make the most use of the available capacity.

Zebi's capacity enhancement features include:

- On-the-fly variable block de-duplication to reduce capacity requirements.
- On-the-fly data compression to help ensure that data is stored in its most dense and space-efficient format.
- Thin provisioning capabilities that allow applications to over-provision capacity while only using what is required.
- High capacity enterprise-class hard disk drives.
- Superior snapshot and cloning capabilities that allow administrators to make copies for backup, restores, data analysis, and mass provisioning of virtual machines without consuming additional disk space.
- Daisy chain scalability that allows administrators to add capacity without impacting performance. Zebi is built in configurations from 10 TB to 90 TB of USABLE capacities and is expandable by JBOD shelf options.

The Tegile Zebi appliance supports in-line compression and block level de-duplication of all data before it is written to disk. It is not a post-processing batch job that is done after all the raw data is written to disk. By combining its unique technology and using high performance Xeon processors, Tegile ensures that the in-line compression and de-duplication does not add any latency to the Zebi.

For backup, organizations typically store multiple copies of their full dataset, which is backed up on weekends. These full copies inherently have a lot of duplicate data. The Tegile Zebi array is able to detect duplicate blocks of data and avoid writing these duplicate blocks of data to disk. This not only saves disk space but also improves throughput speeds because the data does not have to be written again to disk. Tegile Zebi is able to offer similar benefits for other use cases such as virtual server / desktop environments where there is a lot of duplicate data.



Customer example: 2.74 TB is compressed ~60% to 1.47 TB

"I/O wise it has taken everything I have thrown at it without a hiccup and overall performance has been absolutely fantastic."

– Mike Horwath, 1/12 blog
Searching for storage: Tegile

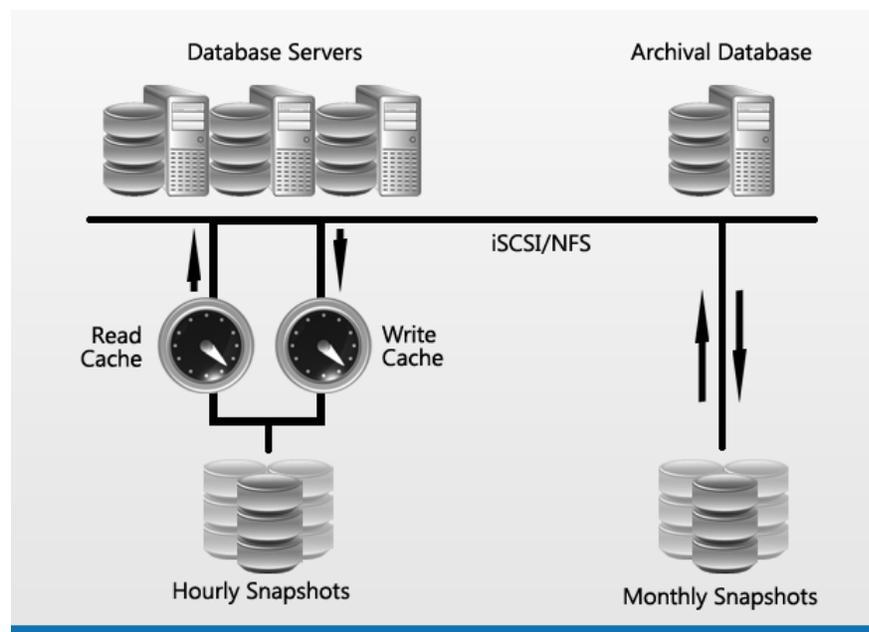
Performance

Zebi is up to 5 times faster than standard arrays allowing organizations to manage higher volumes of simultaneous transactions and additional concurrent users. Zebi incorporates a powerful Xeon processor, large memory, and solid state disk-enhanced read and write optimization to provide high IOPS. It uses line speed SHA256 hash calculations for on-the-fly de-duplication, compression, fast, efficient snapshots and instantaneous restores. This means that in virtualized environments IT managers can easily manage more hypervisors and mitigate the disruptive impact of "boot storms." Additionally, data processing time can be reduced from hours to just minutes. As a result organizations need fewer servers, hard disk drives and software licenses to deliver better services to its IT users.

Zebi's performance enhancement features include:

- Patent-pending Metadata Accelerated Storage System (MASS) technology.
- Intelligent dynamic caching technology.
- High speed DRAM cache.
- High speed solid state disk cache.
- High speed quad core Xeon processors.
- High speed 10Gbps or Fibre Channel connectivity.
- "Prime-time" de-duplication and compression that does not add to latency.
- Non-disruptive snapshots and near-instant restore capabilities.
- Thin replication to reduce network bandwidth requirements by only transmitting changed data.
- Daisy chain performance scalability.

High performance by combining solid state disk flash drives and SATA/SAS drives



"We had racked it up, cabled it up, and they spent their 10 minutes in the data center doing their setup magic. Yep, you read that right; only 10 minutes to set it up."

– Mike Horwath, 1/12 blog
Searching for storage: Tegile

Compatibility/Connectivity

Zebi arrays easily integrate into existing enterprise storage environments without server-side agents so they can work alongside or replace EMC, Dell, IBM, NetApp, HP, Hitachi or other arrays.

Both storage area network (SAN) and network attached storage (NAS) deployments are supported from the same array. Additionally, file-level protocols allow Windows, Mac and Linux applications to share the same storage to create truly unified storage.

Zebi's compatibility/connectivity features include:

- Built-in iSCSI and Fibre Channel block-level SAN protocols.
- Built-in NFS and CIFS file-level NAS protocols.
- 6 built-in 1Gbps Ethernet ports and a 1Gbps IP-KVM "lights out" management port.
- Optional dual-port 10Gbps Gigabit Ethernet .
- Optional dual-port 4 or 8Gbs Fibre Channel.
- Certified VMware compatibility.

Usability

Tegile also built in true unified access into Zebi making it n easy-to-use, yet flexible, platform for virtualization, database and file services. Zebi arrays are fast enough for high transaction Tier-0 applications and real-time Tier-1 applications. They are inexpensive enough, and have the capacity required, for near-line Tier-2 data and Tier-3 backup data.

The easy to use configuration wizard and tiered access controls allows administrators, or designated users, to optimize virtual machines with just a few clicks so they can easily deploy many hypervisors or shares in minutes, not hours. Additionally, the graphs and customizable monitoring worksheets make it easy to identify trends and issues for better planning and efficient optimization.

Zebi can be auto-configured to optimize memory, solid state caches and disk storage by application. For example, as a backup store, Zebi optimizes for high network throughput and in-line de-duplication and compression. As a virtual store Zebi can be auto-configured for high-IOPS, snapshots and instantaneous restores. Administrators can always tweak or create new application storage templates in expert configuration mode.

Zebi's usability features include:

- Browser-based graphical dashboard with tiered access controls.
- Application-specific configuration wizards that allow storage administrators to easily optimize the array without special training.
- Customizable monitoring worksheets.
- Integrated data reduction and data backup features.
- Daisy chain scalability.
- High performance DRAM and Flash-based solid state disk caching and tiering for Tier-0 and Tier-1 applications.



Zebi's informative dashboard

Automatic Use Case Optimization

- High capacity and low costs for Tier-2 and Tier-3 applications.
- Certified VMware compatibility.
- Virtual Machine software agents are not required.

ISCSI LUN under zebipool > Backup

Create a volume/LUN that can be used by services like iSCSI, FC, etc.

Name and Space configuration

Volume Name:

Volume Size:

Thin Provisioning

Initiator and Target Binding: All Select Later All instances in the System

Purpose

Purpose of Usage: Backup Virtualization Database Access Storage

Block Size:

Reliability

Zebi arrays are architected with no single point of failure, rich RAID options and automatic fail-over capabilities. Additionally, Tegile can provide 24x7 support, free software updates and optional on-site support.

Zebi's reliability features include:

- No single point of failure architecture that includes dual hot-swappable controllers, dual power supplies and hot disk spares.
- Automatic error detection and correction capabilities.
- Rich RAID options including Mirror, RAID 5, RAID 6 and Triple-parity RAID.
- Data stored on hard disk drives (Flash drives can wear out in enterprise environments).
- Modular hardware that future-proofs capacity and performance without disruptive down-time.
- 90-day 24x7 support by phone or email.
- Next business day replacement parts.
- Free software updates.
- Optional 4-hour on-site support with on-site hardware kit.
- Optional next business day on-site technical support.

Data Protection

Zebi arrays keep an organization's applications online and protect better than standard arrays allowing IT organizations to improve their Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs).

They come complete with unlimited automatic snapshot and replication features. Automatic snapshots can be scheduled differently for each LUN or file share so administrators can elect to back up critical machines more frequently than less critical ones to save space and improve performance.

Backups are significantly faster and I feel very comfortable with our ability to restore data very quickly from snapshots on disk."

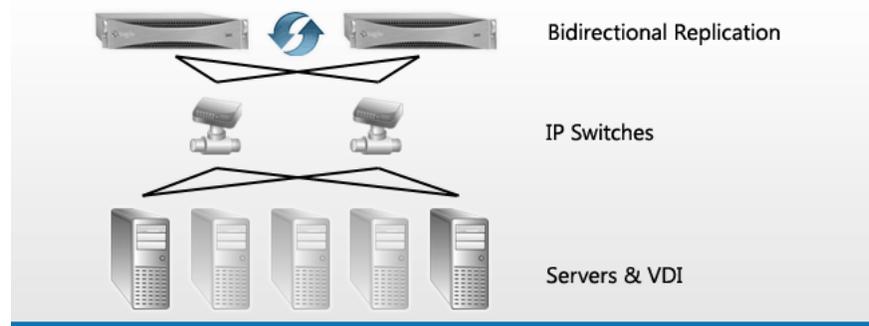
– Arthur Leaser, IT Director, LBBS Law

Zebi's instant restore capabilities allow IT managers to immediately roll back one machine or all machines to a previous state should the need arise. Additionally, the Zebi's automated thin replication feature backs-up only the data that's changed, so even remote backups require less network bandwidth, hardware and administration.

Zebi's data protection features include:

- Automatic non-disruptive snapshot and cloning capabilities.
- One-touch near-instant restore capabilities.
- Automatic Thin Replication feature that only transports what's changed for better performance.
- Automatic fail-over capability where the secondary array can also be used for primary storage.
- Unidirectional, bidirectional, 1-to-n and n-to-1 replication support for maximum flexibility.

Zebi offers Thin Replication and fail-over capabilities



"Tegile's Zebi gives us all the features we need at a much lower price-point and makes it possible for us to meet users' needs while staying within our budget!"

– Jef McCreery, Director of Core Systems, W&L University

Value for Money

With all these features a budget conscious CIO could be forgiven for thinking that Tegile's arrays are just as expensive as arrays from EMC, Dell, IBM, NetApp, HP, Hitachi or other major manufacturer. However, Zebi arrays are priced as low as 1/5th to 1/8 the incumbents.

Zebi's value features include:

- Hybrid DRAM/Flash/hard disk architecture delivers both high performance and high capacity.
- Built-in data reduction technology so usable capacity is greater than raw capacity.
- No additional licenses required for data backup features.

5

SUMMARY

Tegile has taken a fresh look at the seven key requirements for storage today and at how they can best be balanced to support the IT and business needs of the modern data center. One of the key inhibitors to achieving harmony amongst these key requirements has been the handling of metadata. So, Tegile developed its unique, patent-pending MASS technology. By integrating MASS into a new generation hybrid storage system with both solid state disks and hard disk drives, Tegile is able to deliver high-capacity, high-performance storage with all of the functionality needed by the modern data center at a price that is extremely competitive.

Zebi is distinguished by its:



Simplified Architecture

- Hybrid NAS/SAN storage appliance .
- Optimized software and hardware.
- Single license: software, features & data services.

High Performance

- High IOPS for virtualization.
- High throughput for data backup.
- Up to 70% storage reduction.

Economical Pricing

- Up to 80% cost savings over arrays from EMC, Dell, IBM, NetApp, HP, Hitachi and other manufacturers.

With Zebi, Tegile is successfully revolutionizing how organizations in financial, transportation, manufacturing, educational and other industries are overcoming the challenges databases, file sharing, virtualization and backup are placing on their budget-strapped yet rapidly-expanding storage infrastructures.

Specifications may change without notice.

The words Tegile and Zebi are trademarks of Tegile, Inc.
All other marks belong to their respective owners. © Tegile
Systems, Inc. 2012 | Rev. 012612



Tegile Systems, Inc.
Web: tegile.com
Tel.: (855) 483-4453