Tiered Storage Takes Center Stage

Fred Moore, President
Horison, Inc.
www.horison.com

Introduction

The selection of storage technologies has never been greater. Today's choices range from ultra-high capacity, low cost, low performance storage at one end to highly advanced data management functionality and very high levels of performance at the other. Future technology developments promise to further accentuate the differences between the available solutions while increasing the number of choices. Many of these options represent unique levels or tiers of the storage hierarchy. The foundations of tiered storage had their beginnings over 20 years ago when disk, automated tape libraries and advanced policy-based data management software combined to effectively migrate less active data to less expensive storage devices.

At the highest level, tiered storage refers to an infrastructure capable of optimally aligning storage service levels with application requirements. The business case for implementing tiered storage is good and becomes more compelling as the storage pools get larger. However, tiered storage implementations need hardware and storage management software to work together in order to truly provide a seamless active operation and for customers to realize the substantial TCO and ROI benefits.

The example used in the Storage Economics section of his report illustrates this point by comparing the acquisition costs of various single and multi-tiered tiered storage implementations. Using 100 TB of data in the example, the acquisition price of a two-tiered, all disk tiered implementation is 2.1 times more expensive at $765,000 compared to three-tiered implementation using tape for tier 3 archival data costing $359,250. The economic advantages increase further as the storage pool gets larger. In addition, tiered storage significantly reduces energy consumption by allocating data on more energy efficient technologies reducing the TCO (Total Cost of Ownership). Tiered storage offers compelling benefits and they should be strongly considered.

Key point: Tiered storage, as a storage initiative, has been proven to significantly reduce acquisition and operating costs for medium and large scale storage enterprises.

Tiered Storage – Understanding the Storage Tiers

Tiered storage is a data storage environment consisting of two or more kinds of storage, delineated by differences in four primary attributes: price, performance, capacity and functionality. As storage environments grow, an automated tiered storage environment is becoming a necessary architecture, since 1) manual data movement is time consuming 2) the amount of digital data is continually increasing and 3) limited staff resources leave storage administrators stretched too thin. Tiered storage, when properly designed and implemented, can be managed at a much lower cost and as effectively as a single-tiered architecture by deploying tools that can move, migrate, and respond to lifecycle needs of the data or application. To implement a data lifecycle management strategy from a technology perspective, the de-facto standard three-tiered storage hierarchy model of storage tier 1, 2, and 3 is used as the preferred and optimal choice. A fourth tier called Tier 0 based on Flash Memory is now emerging for high performance applications.

The major storage vendors and many of the newer storage entrants have either announced plans for or offer a variety of tiered storage solutions. Very few vendors actually offer a complete tiered storage portfolio including high-performance SSDs (Solid State Disk), RAID arrays and archival tape libraries. Many vendors' tiered offerings are in reality a "disk only" strategy as they only include variations in RPM speeds and price ranges for disk products. Though this is a popular tiered storage approach with most storage vendors, it doesn't serve tier 3 data cost effectively as it forces archival, lower activity data to reside on continually spinning disks. Data that isn't used shouldn't consume energy. Note that a business doesn't necessarily need to use every available tier but the larger the storage pool, the greater the benefit of tiered storage.

A few tiered storage appliances are now appearing that combine disk and possibly flash memory tiers in a single frame. These are usually targeted at a specific application. Presently, businesses must build most tiered storage implementations by integrating a variety of vendors' hardware and software products together. Achieving a unified tiered implementation with multiple vendors' storage devices and management software can be a challenging integration effort, but it can yield significant and lasting efficiencies. Vendors who facilitate this process best will gain market share.

The Tiered Storage Model chart below describes the key attributes of the four tiers in terms of criticality, availability, performance, and technology. A new tier based on solid state disk using Flash Memory named tier 0 is just beginning to emerge and is defined for data that has the highest response time and performance requirements.

Key point: Tiered storage allows an organization to optimize its data storage infrastructure using a combination of storage solutions to lower costs, increase performance and scale technology to address growing storage demands.

Tiered Storage Model

Data tier	**Tier 0**	**Tier 1**	**Tier 2**	**Tier 3**
Amount of data in tier (typical ranges)	1-3%	12-20%	20-25%	43-60%
Primary technology	SSD (Flash)	High - performance disk arrays, FC	Midrange disk arrays, SAS, SATA	Tape libraries, offsite data vaults
Data Classification Category	I/O intensive, response-time critical	Mission-critical, OLTP, revenue generating applications	Vital, sensitive, business important applications	Archives, fixed content, compliance, reference data
Average TB Managed per administrator (open systems)	TBD	30 TB	30 - 100 TB	x TB –> y PB
Availability	99.999%	99.999%	99.99%	99.0-99.9%
Acceptable downtime	None	None	<5 hours/ year	< 1 day/year
Problem response	< 2 hours	< 2 hours	< 5 hours	<24 hours
I/O performance capability	Highest > 1 million IOPs	High 200-300 IOPs	Moderate 100-200 IOPs	Moderate, low
Backup RPO	< 4 hours	< 4 hours	< 12 hours	1 day or more
Application RTO	< 1-2 hours	< 1-2 hours	< 5 hours	<24 hours
Disaster protection	Required	Required	Select applications	Select archives
Data recovery	Mirrored, replication	Mirrored, replication	Scheduled backups	Local and remote backup
Interface connect (open systems)	FC multi-path	FC, multi-path	FC and IP	FC and IP
Power consumption per GB	Low	Highest	High	Lowest

Source: Horison Inc.

Tiered Storage Model – Align Your Data with the Best Tier

If data classification software is not readily available, the template above can serve as a simple starting point to begin the classification process by mapping the required data characteristics for a given file, data set or application to determine the optimal tier. Several observable storage profiles have evolved that help to understand data behavior over its lifespan:

- the probability of reuse of most data declines as the data ages
- the value of data can change over time based on a variety of circumstances
- the amount of data is increasing as it ages as more data is being kept for much longer periods of time than ever before pushing older data into tier 3 status

It is for these reasons that managing and protecting data throughout its lifetime is today's most critical storage management task.

Tier 0: High-performance storage - Tier 0 is just emerging out of tier 1 storage for extremely high performance, high value information that needs to be captured, analyzed and presented at the highest possible speed. The technology solution for tier 0 is SSD, using either DRAM or Flash memory, although new hybrid technologies may soon compete in this space. SSD offers the highest levels of I/O device performance and is the most expensive tier on a $/GB basis and the most cost effective on a IOPs/$ basis. The typical ASP (Average Selling Price) range for Tier 0 is $50-100 per GB.

Tier 1: Primary storage - Tier 1 mission critical data uses the latest-generation of Fibre Channel disk systems and requires high performance, high availability with near-zero downtime and fast RTOs (Recovery Time Objectives) to support customer-facing and revenue-generating applications. These systems carry a premium $/GB price, but this cost is normally justified because slower performance or lower availability systems would directly impact customer satisfaction, business revenues and possibly corporate viability. The typical ASP range for Tier 1 is $7-20 per gigabyte.

Tier 2: Secondary storage - This tier includes midrange disk storage normally with less functionality and at lower price points than tier 1 disk. Tier 2 supports a broad range of major business applications including databases, backup, email, file systems and Enterprise Resource Planning (ERP). Tier 2 solutions must securely store active business data where sub-second response is not necessarily a requirement but reasonably fast response still is needed. Choosing the optimal tier 2 technology is normally a balance between cost, performance and availability. The typical ASP range for tier 2 is $1-8 per gigabyte.

Tier 3: Long-term storage - Tier 3 storage represents the fastest growing segment of the storage industry at nearly 60% compounded annually and tier 3 is the primary future growth area for tape. Tape is the most cost-effective destination for fixed content, compliance and archive data given that tape has a lower price per gigabyte and much low energy consumption per gigabyte than disk storage. Tape has offered WORM (Write-Once-Read-Many) capabilities for several years and has led the way among all storage devices in implementing encryption for data at rest enhancing its position as a storage medium for compliance, legal and any types of sensitive files. Archives retain data that have infrequent access patterns, yet must be protected and kept for prolonged periods of time and many archives are now being kept forever. Data retrieval can pose challenges for archival storage systems, which must provide the capacity and reliability to retain data for many years, protect that data against unauthorized changes and infrequently locate data from the archive on demand in an acceptable timeframe. The typical ASP range for Tier 3 is $.20 to $2.00 per gigabyte and is a function of the ratio of cartridges to drives in a tape library and assumes a 2:1 compression factor.

Tier 3 is characterized by high capacities, moderate to low activity, and long-term if not infinite retention requirements. As most data ages, access activity drops off rapidly and the data moves to an archival status. Without a tiered strategy, the data stays on more expensive tier 1 and tier 2 disk based systems. Tier 3 storage is often labeled as green storage as it has the lowest energy consumption of any tier enabling the elusive IT goal that **"data that isn't accessed shouldn't consume energy"** to become a reality. This is one more reason that tape remains the most cost-effective physical medium for archival data storage and future tape developments indicate this trend will continue.

Data Classification – Understanding Your Data

With or without sophisticated tools, most organizations should be able to classify many of their key applications and have a clear understanding of the value of that data to the business. This pattern of matching data and devices can then assign data to the optimal tier. In practice, however, tiered storage implementations have often been limited due to:

1. Lack of automated data classification and limited deployment of migration tools
2. Lack of support for heterogeneous server and/or storage environments.

Most digital data begins its life on disk and then migrates through the tiers based on a variety of user-defined policies as it ages. In practice, data is classified into four categories and these categories are mapped to the storage tier best fitting those data characteristics.

Data Classification Categories – Understanding What They Mean

I/O intensive data represents up to 3% of online data and includes a variety of response time critical applications. This data typically resides on tier 0 Flash or tier 1 high-performance disk arrays. Application examples include:

- High performance operating system files
- HPC applications
- Select high-performance databases
- Indices, logs, roll files, directories, catalogues
- Database acceleration
- Any data that demands the highest levels of I/O performance

Mission-critical data generally represents 12-20 % of online data and is the most valuable data to an enterprise, used in the critical business processes, revenue generating or customer facing applications. Data recovery times, the RTO (Recovery Time Objective), ranging from a minute or less up to an hour or two, are typically acceptable. Mission critical data usually reside on tier 0 or tier 1 storage based on the performance requirements. Application examples include:

- Oracle/MySQL and mission critical databases
- Reservation systems
- ATM, Point of Sales
- Virtual machines
- E-mail (in some businesses)
- Online transaction processing (OLTP)

Vital and sensitive data typically averages up to 20-25 % of all data stored online and represents the majority of data stored on disk. However, vital and sensitive data doesn't require instantaneous recovery for the business to remain in operation. Data recovery times can normally take up to four or five hours without causing major operational impact. The wide range of vital and sensitive data applications include:

- Business-critical applications
- Oracle/MySQL data bases
- Web servers and applications
- Data protection – backup, recovery data, security systems
- Image capture and retrieval
- Application development and test
- Transactional workloads
- Data warehousing, ERP, reference data
- Cloud storage

Archival data presently represents approximately 43-60 % of all data stored online making it the largest category and at >60% CAGR, is the fastest growing data classification segment. Low activity, long term and infinite retention periods for data such as e-mail archives, compliance and legal records, medical information and images, video archives, photos and fixed content all fit this profile. Lost, corrupted or damaged data can be reconstructed with minimal effort, and acceptable recovery times can range from hours to days. Archival data is typically stored on the lowest cost storage systems, primarily automated tape storage, or other removable storage systems, in local and/or remote locations. Tier 3 data is the primary market for future tape growth and the prime opportunity for outsourcing data to cloud or other types of service providers. Tier 3 application examples include:

- Long-term data retention, archive, and backup
- E-mail archives
- Compliance, medical records and images, and legal data
- Unstructured, file-based data, documents
- Scientific data
- Video archives, movies, audio, blogs, photos
- Wikis, Collaboration, Social Networks, Web 2.0. SOA and Cloud Storage
- Business continuance
- Video surveillance and security system history and archives
- Off-site media storage, remote data vaults

The deployment of comprehensive automated data classification tools remains limited. As a result, while many data centers have a form of tiered storage but the tiers are not integrated in any meaningful way and function as data islands. Even as CIOs lament the exploding growth of disk farms, which are devouring IT budgets and overcrowding data centers, they continue to maintain expensive tier 1 or tier 2 disks partially full of data having little or no activity for years.

Key point: With or without sophisticated tools, organizations should be able to characterize and stratify data and storage. Most businesses can define a finite number of tiers to classify storage and data.

Storage Economics – Understanding the Implementation Cost

Since not all data, applications, and storage resources are created equal, the amount of management effort applied to storage tiers does not need to be equal either. Lower-tiered applications tend to have relatively fewer labor (management, provisioning, sizing, reporting) efforts assigned. Tier 1 disk storage typically receives the most management attention. For open systems environments, a storage administrator manages about 30 TB of tier 1 disk space on the average, but can normally manage the capacity of an automated tape library frequently reaching several petabytes. This wide management gap continues to provide automated tape with lower management costs per TB and a lower TCO than managing the same amount of data on disk. Therefore, scarce labor resources can be saved and used for managing the more critical areas of the storage infrastructure. The chart below lists ASPs (Average Selling Prices – not list price) for a variety of storage systems. This chart serves as a guideline for planning purposes and comparisons only as prices vary widely based on the configuration, size of the transaction, extra features, account criticality, and other economic considerations.

Average Selling Prices (ASP) for the Storage Hierarchy

Subsystem Category	ASP Range ($/GB)	Notes
Solid State Disk (DRAM) FC, SCSI	$300-500 (avg. $400/GB)	ASP range based on capacity and high-availability features
(Flash) SSD/HDD	$50-100 (avg. $75/GB)	ASP based on single Flash SSD
Enterprise disk SCSI, FC	$7-20 (avg. $13.50/GB)	Includes controller, cache & drives. Add-on storage modules with no controller and cache are lower.
Midrange disk SCSI, FC, SAS	$1-8 (avg. $4.50/GB)	Includes controller, cache & drives. Add-on storage modules are lower.
Optical disc library (Blu-laser) SCSI/WORM	$5-20	Includes drives, media and library (No longer a data center technology)
Economy disk SATA, JBOD	<$1-5	Price range varies widely based on capacity and array configuration.
Integrated Virtual Tape Library (tape library with disk buffer)	$.25-$2.00 (non mainframe) <$.20 (mainframe, varies by configuration) avg. ($.20/GB)	Includes tape drives, media and library and uses 2:1 compression to determine library capacity

Source: Horison Information Strategies (July 2010) Average Selling Prices used, prices vary by configuration and vendor.

The example below compares the initial acquisition cost of 100 TB of storage for a disk-only tiered strategy to several tiered storage implementations. The ASPs are calculated from the table above using the mid points of each range as appropriate and serve as an example. Customers can utilize their own respective storage cost metrics when performing a similar analysis. Please use as input actual price quotes from your storage vendor to determine costs for any specific proposal. The average amount of data in each tier uses the industry averages from the Data Classification Model described earlier. For 100 TB of data, the acquisition price of a two-tiered, all disk tiered implementation is 2.1 times more expensive at $765,000 compared to an active three-tiered implementation using tape for tier 3 archival data costing $359, 250.

Economic Considerations for a Tiered Storage Implementation (Example)

Total capacity	**100 TB**				
Storage Tiers	**Tier 0**	**Tier 1**	**Tier 2**	**Tier 3**	**Total**
ASP/GB by tier (range avg.)	$75.00	$13.50	$4.50	$.20	
Total price per tier (if all data on one tier)	$7,500,000[1]	$1,350,000	$450,000	$20,000[1]	
Avg. data allocation % by tier	2% (optional)	15% (13% if tier 0)	32.5%	52.5%	100%
Data allocation after classification process	2 TB (optional)	15 TB (13 TB if tier 0)	32.5 TB	52.5 TB	100 TB
Total price 2 tiers (disk only tiers)	----	$472,500 (35 TB data)	$292,500 (65 TB)	----	$765,000
Total price 3 tiers (disk and tape)	----	$202,500	$146,250	$10,500	$359,250[2]
Total price 4 tiers Add tier 0 (SSD)[3]	$150,000	$175,500	$146,250	$10,500	$482,250[4]

1 - For comparison purposes only, not an actual storage solution
2- Sum of tier 1, 2, 3 prices representing typical tiered hierarchy
3- Re-allocates 2 TB of high performance data from tier 1 to tier 0
4- Sum of tier 0, 1, 2, 3 prices

Key point: Tiered storage reduces the overall purchase price, monthly license fees and maintenance costs and enables storage administrator efforts to be focused on more critical tasks. The economical benefits of tiered storage increase as the storage pool gets larger.

Oracle and Active Tiered Storage

Few companies offer a complete tiered storage solution that includes a full selection of storage tiers along with effective HSM type of software. With a range of disk offerings, an industry leading family of StorageTek automated tape solutions and key HSM software from its recent acquisition of Sun Microsystems, Oracle has positioned itself to become a leading provider of active tiered storage.

Critical to Oracle's tiered storage strategy is its HSM software called Sun Storage Archive Manager (SAM) with its Sun QFS software (Quick File System). QFS is an open source, POSIX compliant file system. Together, SAM-QFS provides policy-based tiered storage management and shared file access in Open Systems environments. It is this software that makes Oracle's tiered storage become an "active" tiered storage solution. Without policy-based software to manage the storage tiers, the tiers become "islands" of labor-intensive storage to manage and some of the up-front cost savings can be lost to increased management costs. SAM-QFS can manage from gigabytes to 4 petabytes of disk tiers backed by 10s of petabytes of tape in tier 3 for archival storage. Files are copied to archival storage in the background, and dynamically and transparently retrieved whenever accessed. SAM-QFS reduces capital and

operating costs through policy driven, dynamic (active) data movement across tiers of storage, local and remote, based on data classification, and transparently migrates data to the most cost effective storage tiers, an important ingredient for tiered storage.

Moving data to newer technologies can be time very consuming and is often performed as a manual task in many businesses. An additional benefit of an automated tiered storage environment is that it can relieve much of this burden. In addition to moving data through tiers of storage based on its classification, SAM can also move data from older technology onto new technology and retire the old devices or media non-disruptively automating much of what has historically been a manual process. As today's longer data retention requirements outlive technology, it is important to the sustainability of the data to have it always on a device and in a format that can be read at anytime.

The implementation of Storage Archive Manager and QFS has proven to be an effective tool for developing archiving practices for specific industries or applications. When an Oracle integrated storage stack of software, hardware, and services are combined using SAM-QFS, the foundation for highly cost effective and scalable tiered storage architecture can be built. With the acquisition of Sun Microsystems, Oracle becomes one of the few companies to deliver a true active tiered storage solution by providing a broad range of hardware coupled with robust HSM type software.

Oracle Tiered Storage Positioning

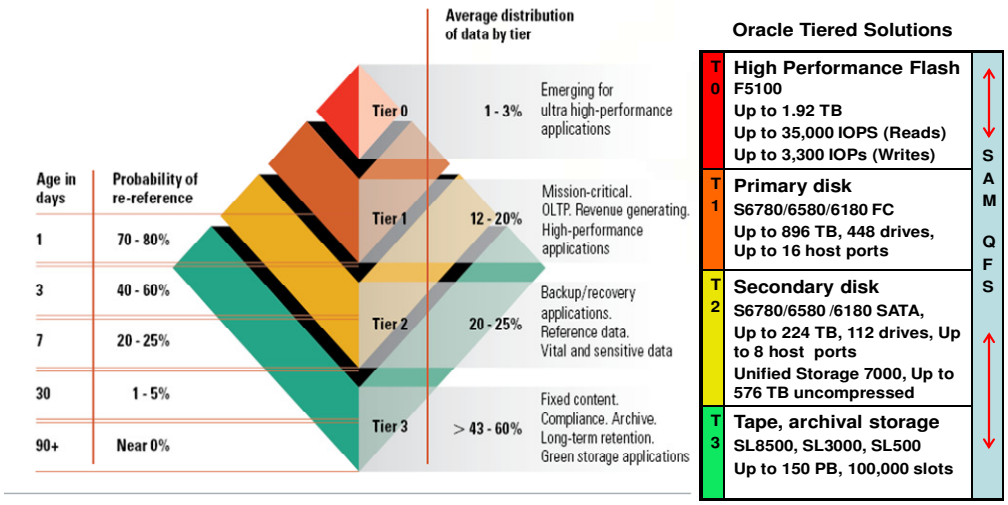

Comment [OU1]: As of last week, SL8500 has now 150 PB (LTO 5), also to make things simpler, suggest we take out the drives and just show the libraries

Comment [GF2]: Do we want to put VSM in Tier 3? I don't think so. For the purpose of this paper I would just leave it out. Too difficult to explain as to whether it is T1 or T2.
MS: Agree take VSM out

Key point: Oracle becomes one of the few companies to deliver a true active and automated tiered storage solution by providing a comprehensive range of hardware coupled with robust HSM software.

Conclusion

A multi-tiered storage system with automated data management provides the optimal solution for managing the 21st century data explosion. As noted in the previous example, for 100 TB of data, the acquisition price of a two-tiered, all disk tiered implementation is 2.1 times more expensive at $765,000 compared to an active three-tiered implementation using tape for tier 3 archival data costing $359, 250. The economic advantages for a tiered storage implementation increase further as the storage pool gets larger. Investment protection increases by using storage tiers such as those offered by Oracle that scales in both capacity and performance as the overall storage requirements increase. This avoids vendor lock-in along with frequent and costly replacement of older technologies. Organizations that do not archive their data take a large risk if they are ever involved in a civil or criminal action and are required to produce historic documents. Assigning monetary value to the business value of data is a concept that businesses don't always have a grasp of although it would seem to make a great deal of sense to do. The business value of data can be described in terms of lost revenue if specific data is not available to an application. Various surveys indicate that the loss of mission critical data can now reach or exceed $5 million per hour depending on the market segment. However, organizations that do understand the value in their information have found it to be a business advantage and become a leader in their industry as they are able to achieve higher availability and consistent service levels and minimize any type of business outage. Overall, active tiered storage provides the best solution for supporting service level agreements at the minimal cost and with highest efficiency. Pricing levels vary by tier and the larger the amount of storage, the greater the overall financial benefits from a tiered storage implementation.

With a growing number of cost-effective storage solutions in each tier today, it makes sense to undertake the process of classifying data in order to optimize storage infrastructure expenses and the more storage, the greater the benefits from this process. The data classification process can be basic or use advanced software to determine the value of data or an application and to identify characteristics of the optimal type of storage subsystem to meet those requirements. As a result, an effective tiered storage model flows from the data classification process and these processes are linked. Mainframe users have been effectively addressing many of these issues for years. For the vast majority of businesses however, they have yet to tackle this issue in some form before the growing amount of data becomes completely unmanageable. If you are one of those businesses, an ***active*** tiered storage implementation should be a priority.