



Tiered Storage Saves Money — Getting The Most Out Of Your Storage Infrastructure

Analyst: Michael Fisch

Management Summary

When boarding an airplane, most of us walk past the passengers in first and business classes and take a seat with the huddled masses in coach. Not that we prefer close quarters and salty pretzels instead of lots of legroom, decent food, and free drinks – but all of the extra benefits costs more. **Those who want inexpensive transportation fly coach; those who can afford to travel in style fly first class; and in between, there is business class.**

Tiered storage makes a similar proposition. Like passengers on an airplane, not all information requires or can afford the same level of service. Some requires high performance and availability, while other information will find a lower quality of service (QoS) to be adequate – at a correspondingly lower cost. **The idea of tiered storage is to segregate data into classes with similar requirements and create storage pools with the appropriate levels of performance, availability, recoverability, security, and cost in which to contain them.** The net result is a significantly lower overall cost of storage, without sacrificing the needs of business applications.

This is not a new concept; in fact, tiered storage has existed for quite some time in the mainframe world. **The new part is the great need to do it in open systems environments, coupled with storage advances that now make it practical.** Ever-increasing stockpiles of corporate data and relatively flat storage budgets are driving the pursuit of greater cost efficiencies. Recent advances in storage networking, storage virtualization, and data management make it much simpler to create and manage multiple tiers. In short, the time has come for tiered storage.

The beauty of tiered storage is that it can lower storage acquisition costs today and every time additional capacity needs to be purchased. **Rather than waiting for the ethereal long-term savings and intangible returns of other IT projects, tiered storage can deliver an immediate return on investment.** This is something CFOs like to hear, especially nowadays.

So, if you are considering a tiered storage solution, the good news is that there are a wide variety of technologies and techniques available. This includes various storage types and features as well as management and automation tools. **To help sort it out, read on for a summary look at the benefits, considerations, techniques, and technologies for implementing tiered storage.**

IN THIS ISSUE

➤ The Need for Tiered Storage.....	2
➤ More Cost-Effective Procurement.....	2
➤ Creating a Tiered Storage Infrastructure	4
➤ Conclusion	6

The Need for Tiered Storage

Unlike people, all information is *not* created equal. Some information needs to be accessed very quickly; other does not. Some information requires fast recovery from faults and disasters; other does not. And so on, and so forth. These distinctions are not merely academic; they provide an opportunity to segregate data into different storage pools – and like passengers on an airplane, not all data needs to fly first class. An enterprise can *save money* by assigning critical information to robust, high-end storage and less-critical information to more basic and less-costly storage. **This is the essence of tiered storage – to establish multiple pools of storage with different qualities of service (QoS) in order to optimize storage expenditures.**

The idea of tiered storage is not new. In fact, it has been around for many years in the mainframe world under names like hierarchical storage management (HSM) and archiving. What is new are technologies that make it more practical and doable in open systems environments, such as storage area networks (SANs), network-attached storage (NAS), storage virtualization, and data management.¹ **Coupled with the perennial problem of expanding storage capacity requirements and costs, tiered storage will take on new significance and urgency as a way to keep costs under control.**

Storage QoS can be described along several dimensions: performance, availability, recoverability, security, retention periods, and cost (*see the box on the following page*). Generally speaking, higher levels of service cost more. For instance, high-end RAID arrays with Fibre Channel drives are much more robust than JBOD², but at a correspondingly higher total cost of ownership (TCO). This tradeoff between cost on one hand and performance, availability, and so forth on the other is the basic distinction between each tier of storage. Enterprises can cover several points along the spectrum by creating multiple tiers of storage.

The key is to match processes and data to the appropriate QoS and tier of storage. Factors that determine the QoS of a file or set of files (e.g., a database) include the criticality of the applications that use it (i.e., their relative utility or importance) and the importance of the user. For instance, an application that performs online transaction processing typically requires high performance and

availability because it supports ongoing business operations. Slow transactions can lower business productivity, and downtime might cost many thousands of dollars per minute in lost revenue, not to mention the impact on customers. This data should be assigned to the highest storage tier. However, older or infrequently-accessed records may not require fast access and could be moved to slower, less-costly storage. Thus, a single application might access data in more than one tier. In another example, a tool for data mining uses separate, point-in-time copies of production data for its analyses. Though the tool benefits from good performance, it does not require high availability because the data is only used temporarily then re-created for the next analysis. A storage tier with decent performance, less-than-highest availability, and lower cost would be optimal. In short, the characteristics of the applications and relative importance of the data and user will usually determine the required tier of storage.

The result of tiered storage is an infrastructure that minimizes costs without sacrificing the needs of business applications. **By making intelligent distinctions between classes of information and its use, an enterprise can optimize one of the most critical and expensive components of its IT infrastructure – storage.**

More Cost-Effective Procurement

More specifically, tiered storage benefits enterprises in two ways.

- **Freeing up high-end storage and delivering an immediate payback.**

Large enterprises will generally find their storage is weighted toward the high end. After years of being sold “only the best” by storage vendors, and having assumed all information requires similar QoS, enterprises may find they can migrate some data to less-costly storage and free up space at the high end. As an illustration, assume an enterprise called XYZ Company has 10 TBs of data residing on high-end storage arrays, but only 5 TBs actually require very high performance and availability. It could migrate the other 5 TBs of less-critical data to mid-range RAID arrays, maybe with SATA drives, which deliver less performance and availability but cost, say, one-sixth the price of high-end storage. By deploying \$20,000 of mid-range SATA storage for the 5 TBs, XYZ Company frees up \$120,000 of high-end storage, for a net immediate payback of \$100,000 and an effective ROI of 500%³ – not too bad. The free space at the high end will be needed

¹ See *Networked Storage – Evaluation Criteria* in...

² *Just a Bunch Of Disks* – JBOD are simple disk arrays that lack an intelligent storage controller for performing RAID, caching, clustering, and other advanced storage features.

³ This figure does not include any costs associated with installing or administering the new storage or data migration.

Storage QoS – Performance, Availability, and Cost

While not an exhaustive list, three fundamental dimensions of storage QoS are performance, availability, and cost.

Performance describes the speed of storage access. It answers the question: “How quickly can I access or store the information?” There are two ways to look at it. First, *bandwidth*, which is akin to the width of a pipe, describes how much data can be moved per unit of time, usually in terms of MB/s. It is particularly relevant when moving large, sequential blocks of data like audio or video files. Second, *throughput*, usually expressed in inputs/outputs (I/Os) per second (IOPS), measures the speed of processing many, small I/Os. It is an especially important metric for transaction-processing applications. Throughput is directly related to *access time* – the time it takes a packet of data to move from application to storage or vice versa. Access time is like the length of a hose because it affects how long it takes water to come out from when the valve is turned on (assuming the hose starts out empty). There is also a third aspect of performance that describes how long it takes to retrieve data that has been archived in nearline or offline storage, such as in a backup or archiving system.

Availability is the percentage of time that storage is accessible (e.g., 99.999%). It answers the question: “How reliably can the information be accessed?” Most significantly, availability is described by the three Rs of redundancy, remoteness, and recoverability. *Redundancy* implies redundant system components for failover purposes and multiple copies of data to guard against loss or corruption. Features like RAID, point-in-time or snapshot copy, mirroring, and even backup enable redundant copies of data. *Remoteness* extends the concept of redundancy by placing systems and data copies at facilities far enough away to be safe from local disasters. Remote mirroring and electronic vaulting are technologies that deliver remoteness. *Recoverability* is the time it takes to recover from an interruption, such as a corrupted database or failed storage array. Availability is also affected by the durability of the storage media (e.g., optical, tape, hard disk). Downtime associated with adding incremental capacity, changing storage configurations, or performing backups can affect availability, as well.

Another important dimension of storage QoS is cost or, more specifically, the total cost of ownership (TCO). **Storage TCO includes both acquisition cost as well as operating cost, which reflects the long-term costs of managing storage, especially in terms of the skilled labor required.** Operating cost is the more significant of the two since it can amount to several times the cost of storage acquisition, over its useful life.

in due time since enterprise data grows at rates of 30 to 100% per year or more.

- **Lowering the cost of all future storage purchases.**

Continuing with the previous example, XYZ Company can defer purchasing high-end capacity until the liberated 5 TBs is filled. Not only does this save the internal cost of capital (e.g., 10% per year) for the length of the deferral, but disk prices also drop about 20% per year, so the price it finally pays for the incremental high-end capacity can be much lower. Moreover, assuming XYZ Company’s data requirements are the same going forward, it can continue to purchase storage on a 50/50 split between high-end and low-cost, thereby cutting the price of future storage purchases by over 40%. **So you see, tiered storage is the gift that keeps on giving.**

There is a tradeoff to consider, though. While tiered storage significantly lowers acquisition costs, it can also affect operating costs by adding a degree of management complexity. Installing and administering additional types of storage imply a learning

curve and incremental management activities. Data migration between tiers of storage is also a factor, both initially and potentially on an ongoing basis. The good news is that there are software and hardware solutions available for simplifying and automating the management of heterogeneous storage devices – and even for tiered-storage environments, in particular. The next section will cover some of these solutions that reduce operating costs.

Creating a Tiered Storage Infrastructure

As in building a house, there are a variety of approaches and architectures for constructing tiered storage. First, there are the building blocks – various storage devices and media as well as advanced software features for improving QoS. Second, there are the tools and methods of management, which is the glue that holds the structure together. Considering all of these options can seem confusing and, unfortunately, there is no one-size-fits-all approach. **The important thing is to know your enterprise’s situation and needs and choose one good solution (among many).** A simple step in the right direction can deliver a great deal of value.

The box on the following page describes various storage types and features that can be used to construct storage tiers. While many enterprises have a high-end or mid-range RAID array as primary storage and tape for backup and archiving, this represents a sort of baseline from which to build a more advanced storage profiling solution. **Ideally, an enterprise should deploy one or more tiers of storage for each category – online, nearline, and offline** - online for the best performance and availability; nearline for less-costly storage that is still automatically accessible; and offline for backup and archival. The net result is an optimum, or at least better, balance of performance, availability, and cost.

Methods of Management

Creating storage tiers is only half the picture. **An enterprise must also consider the three Ms of management: *managing* storage tiers, *migrating* data between them, and *measuring* actual QoS delivered.** Ongoing management ensures storage systems continue to operate properly and data is protected and available. Data migration is highly desirable because greater cost efficiencies can accrue by migrating data from online to nearline and eventually to offline storage as its usefulness declines over time (i.e., lifecycle management). The alternative is to assign data permanently to a storage tier. Measuring QoS is very helpful because it confirms whether the storage infrastructure is meeting QoS objectives for internal or external storage “customers” and, if not, alerts the IT department to take corrective action. However, this feature is in the earlier stages of maturity and deployment.

Like the variety of storage types and features, there are various methods of managing tiered storage.

Manual Management

Manual management implies an enterprise does not use any special software to simplify or automate the management of storage tiers. Administrators must take it upon themselves to manage the different devices in each tier. (Some storage vendors offer device managers for both mid-range and high-end storage arrays in their own product lines that can simplify these tasks.) Data migration between tiers is possible but labor-intensive, if performed on a regular basis. Data copying tools can make the process faster and less disruptive, but it still requires human intervention. A manual approach is most practical when storage tiers are dedicated to specific applications. **While the manual approach can work for simpler tiered-storage scenarios, it does the least to lower operating costs, which are**

largely due to the costs of skilled labor.

Archiving (Automated Data Migration) Software

This software category is primarily concerned with automated data movement or migration. *Archiving* software automatically moves data between storage tiers based on factors like age, usage frequency, and content. It places pointers on primary storage to provide seamless (though slower) access to migrated files and records. For databases, it prunes older records and archives them to keep the size of the database manageable and maintain good performance. **These solutions reduce operating costs by automating data migration based on policies and ensuring moved data is still accessible.** Archiving software will play a key role as more enterprises implement tiered storage.

Storage Virtualization Platform

Storage virtualization platforms are a newer, growing technology that makes it easier to manage multiple, heterogeneous storage arrays and organize them into separate pools or tiers. These products come in a variety of architectures (e.g., in-band, out-of-band) and their basic function is to present multiple, heterogeneous storage arrays to servers as a cohesive and dynamic pool of capacity that can be carved up, expanded, and reassigned as needed. In essence, they provide a layer of abstraction over the physical storage that simplifies management and lowers operating costs. Some also offer advanced data management features, such as point-in-time copy and remote mirroring. Virtualization platforms will be an increasingly important technology, and one of its key benefits is to facilitate tiered storage.

Storage Management Software

Heterogeneous storage management software is another technology that simplifies life for storage administrators. While virtualization handles block/volume aggregation, storage management software operates at a high level to monitor, report, control, and automate diverse storage resources. Like a conductor in front of a symphony orchestra, it attempts to direct the smooth operation of heterogeneous environments by managing storage resources, availability, planning, provisioning, and so forth, from a top-down perspective. It typically leverages open SMI-S standards as well as vendor-specific “hooks” for greater control. **While features vary between products, the bottom line is that they can help minimize the work that administrators must do to manage multiple tiers of storage.**

The Building Blocks: Storage Types and Features

When building storage tiers, enterprises need devices with different QoS characteristics. Fortunately, there is a gamut of technologies from which to choose. Solid state disk is by far the fastest and most costly storage because it uses electronic memory rather than magnetic or optical media. It functions as a performance accelerator by storing a small percentage of frequently-accessed records and files such as database transaction logs, indices, and temporary tables. High-end global cache RAID arrays and mid-range RAID arrays with Fibre Channel or SCSI disks function well as primary storage – with high-end arrays offering greater performance, availability, and cost for more-demanding applications. Recently, a number of storage products based on less-costly and less-robust ATA/IDE drives have arrived on the market, including mid-range arrays for primary storage as well as platforms for secondary storage, backup, and archival. JBOD, which lacks the redundancy and reliability of RAID, can be a good fit for less-critical data. Optical storage (i.e., DVD, CD, MO) performs well as a nearline archival solution because it is relatively inexpensive, provides access within seconds using a library or jukebox, retains data for long periods, and is removable. Finally, tape is frequently used for backup and archival because it is the least-expensive media, retains data for long periods, and is also removable. Tape is also the slowest media because data is accessed sequentially, not randomly, as with the other media, but using a virtual tape technology can mitigate this.

There are also a number of advanced data management features for improving storage QoS. RAID, for instance, encompasses a variety of techniques for protecting data from drive failures and increasing performance by striping data across multiple drives. Cache memory on storage controllers also improves performance. Point-in-time and snapshot copies can increase availability by allowing non-disruptive tape backup, data warehouse loading, etc. They can also serve as backup data sets for fast recovery, if production data becomes corrupted. Remote mirroring allows for business continuance in case of a local disaster. Dual pathing between servers and storage not only improves performance but also provides an alternate path in case of failure. **Enterprises can apply any or all of these features to improve the QoS of a storage tier – keeping in mind there is a financial cost associated with each one.**

Finally, in the context of tiered storage, devices are often grouped into one of three categories: online, nearline, and offline. *Online* means data can be accessed automatically and instantaneously – such as within milliseconds. *Nearline* means the data can be accessed automatically but at a slower speed, such as a second or more. *Offline* means the data typically requires human intervention (i.e., finding and loading a tape) to access it. Ideally, a tiered storage solution would contain element(s) in each category.

Conclusion

The concept of tiered storage is an idea whose time has come. It is not a new concept, but several market factors have aligned to bring it to the forefront. First, enterprises simply must lower storage costs in the face of spiraling storage capacity requirements and its increasing portion of the IT budget. Second, the proliferation of storage networks make it possible to create storage pools and more easily migrate data between them. Finally, the emergence of storage virtualization and management technologies greatly simplify the operational aspects of heterogeneous environments that tiered storage implies. (Conversely, the benefits of tiered storage may be a major reason to adopt these technologies.) **In short, the need and means are in place today.**

Tiered storage can significantly reduce storage acquisition costs – both up front and on an ongoing basis. Furthermore, it should be relatively

easy to sell to upper management because it delivers an immediate payback. Executives tend to be more favorable toward projects with swift, certain returns versus hazier long-term benefits. **In the never-ending corporate quest to cut costs, tiered storage appears to be low-hanging fruit ripe for the picking.**



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over ten years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at mike.fisch@clipper.com or at 781-235-0085 Ext. 211. (Please dial "211" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and "clipper.com" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "Navigating Information Technology Horizons", and "teraproductivity" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.