

TECHNOLOGY ASSESSMENT

Tiers of Joy? What's Going on with Tiered Storage?

David Reinsel
Richard L. Villars

Benjamin Woo

IDC OPINION

IT managers are scrambling to keep up with increasing storage requirements within their datacenters while facing shrinking budgets and greater scrutiny of datacenter energy and space costs. A myriad of strategies and solutions, including hardware and software, are emerging to help IT managers address these issues. Effective implementation of tiered storage is an important solution for addressing these challenges. However, IDC finds that the concept of tiered storage, while not new, still is not being leveraged to its full extent, and many times not at all. While a number of elegant point solutions are now on the market, the overall messaging is poor, and worse, the packaging is problematic. Organizations still struggle to find software tools that help them manage information assets across a heterogeneous, tiered architecture. Vendors desire to have stable, relatively long-lived product families, while customers loathe the thought of migrating their storage infrastructure to new, potentially different architectures. Hence, the industry seems to be looking for a technology and requirements plateau upon which to build its next-generation storage solutions. IDC believes the time is ripe for vendors to come to market with tiered storage solutions that meet the requirements that customers have today for their datacenter environments. These solutions must:

- Simplify storage infrastructures as opposed to adding complexity
- Enable far greater automation of data classification and data movement
- Accommodate the rapid pace of storage technology evolution, as well as the increasing diversity of data types

IN THIS STUDY

This IDC study evaluates today's tiered storage environments. It reports on their popularity, their successes, their shortcomings, and the future opportunities for storage vendors.

SITUATION OVERVIEW

Tiered storage is a term that is used to describe a data storage infrastructure (or architecture) that leverages the installation of two or more tiers of storage that can meet multiple business objectives based on characteristics such as cost, acceptable performance, and capacity/density.

The concept of tiered storage is not new. Architectures leveraging various tiers of storage have been used for years. In fact, most end users, especially those with larger datacenters, have sought to establish efficient tiered storage environments (with varying degrees of success). These efforts are a way of improving capital efficiency while still meeting the capacity, performance, security, and regulatory needs of their organizations — within a defined budget.

As information assets expand and diversify, extending the effective use of tiered storage is moving from a desirable attribute to an absolute necessity. In these difficult economic times, customers must deal with unavoidable data growth, and the storage industry remains under considerable evolutionary pressure.

In the past, the most basic tiered storage environment leveraged a tier of primary hard disk drive (HDD) storage that was backed up and archived to a tier of tape storage. This could exist within a single server or in a networked storage environment. Today's storage environments, especially for larger companies, are much more complex and involve a variety of disk-based tiers. In the future, they will include multiple tiers of solid state drives (SSDs) as well.

To date, however, the vast majority of deployed tiered storage architectures are limited to two tiers of disk-based storage — performance-optimized drives and capacity-optimized drives (e.g., 10,000rpm/15,000rpm Fibre Channel [FC], SCSI, and Serial Attached SCSI [SAS] HDDs and 7,200rpm Serial ATA [SATA] HDDs, respectively).

In today's more diverse IT environments, this simplistic approach to tiering is insufficient. Today, IT architects need to consider six major factors:

- ☒ **Performance** — By leveraging techniques such as overprovisioning and short stroking performance-optimized drives, an organization is able to create a tier of storage that provides superlative performance metrics (compared with the rest of the storage infrastructure) to satisfy high I/O applications such as equity trading applications or reservations systems.
- ☒ **Cost** — Conversely, by using (typically) lower-cost capacity-optimized drives, organizations can create a tier of storage that is focused primarily on long-term

archiving of data over time (as an example). Therefore, while the preservation of the data is critical to the organization, instantaneous data request is second to the cost-effective nature of simply having relatively quick access to the data for future needs or requests.

- ☒ **Function** — As the number of regulatory and legislative requirements increases for the preservation of the actual data, as well as integrity and security, organizations may be required to create one or more tiers of storage in which data is not only preserved in its original stored state but also tagged with audit trails to track how the data has been accessed and by whom or by which applications.
- ☒ **Environment** — With unpredictable power, cooling, and real estate costs, an organization may consider creating a tier of storage focused on energy efficiency (e.g., through the use of solid state storage technologies).
- ☒ **Reliability** — Depending on the perceived or real criticality of the data being stored, an organization may also manage its data based on reliability factors. For example, one organization may consider mirrored and replicated as the appropriate strategy for mission-critical data, and the use of RAID 3, 5, or 6 as appropriate for less critical data. On the other hand, another organization may consider an array of disks protected by the properties of RAID 6 suitable for its mission-critical data.
- ☒ **Efficiency** — An additional motivation to move to tiered storage can be the desire for a more efficient storage infrastructure in terms of simplicity of management. Storage consolidation, virtualization, and data management can all individually increase the complexity of a storage infrastructure, something that IT managers wish to avoid. Integrating such features not only improves the efficiency of storage but also can drive up storage utilization.

Considered alone or in combination, these industry dynamics compel IT managers to move toward a storage infrastructure that supports a growing range of storage tiers. Nevertheless, the industry is far from realizing the full promise of tiered storage.

There are a number of successes that can be highlighted, but also a number of significant limitations in terms of manageability and utilization. A number of additional enhancements must come to the market before end customers, no matter what the driver, can fully realize the benefits of tiered storage.

Tiered Storage Firsts

As it became apparent that companies could no longer afford to grow their storage environments on the back of traditional tier 1 (aka performance optimized) storage, lower-cost options were explored. Initially, companies sought to leverage inexpensive SATA drives — the drives designed for desktop PC applications. However, the industry soon realized that these drives were not capable of performing in multidrive environments. The excessive vibration, time-out errors, and excessive thrashing these drives experienced resulted in early failures and nonfunctioning storage systems — something not acceptable within enterprise storage.

HDD vendors were quick to respond by launching enterprise-capable capacity-optimized drives that were leveraged off of their desktop PC HDD platforms but integrated with special hardware and firmware, which enabled the drives to work within the always-on enterprise environments and under more extreme workloads. As the capacity-optimized SATA drives became a viable and sustainable enterprise option for some workloads, both established and emerging storage vendors began to integrate them into their storage portfolios.

The idea was to provide end customers with an option to migrate old, rarely accessed data from their tier 1 storage arrays to lower-cost, lower-performing storage (aka tier 2) but keep that data online and available within milliseconds. In reality, this kind of data aging across multiple applications and data sets is not possible with most of today's storage solutions. Most companies simply tier data by individual applications. Very few storage solution vendors (e.g., Compellent in the block storage world and F5's or Brocade's file virtualization solutions in the file storage world) provide a dynamic and automated migration of data up and down a storage architecture with multiple tiers. Pillar Data is another early implementer of tiered storage, and already has integrated SSDs into its product solution set.

Today, the most widespread use of intelligent tiering is for making copies of data for disk-based backup, test/development, or compliance. Products like EMC's Centera, NetApp's NearStore, and Data Domain's DDX initially were designed to meet this specialized but fast-growing need. By adopting a replication storage tier, end customers were able to not only conserve precious performance tier storage with additional headroom (thereby staving off an additional capital expenditure — at least for the time being) but also offload menial but important and I/O-consuming tasks like backup to tape. This achievement alone helped to instill into the minds of IT managers, as well as system OEMs, the need to develop a richer set of tiered storage options.

One company that looked to extend the range of tiered storage options was COPAN, which brought to market the idea of massive arrays of idle disks (MAID). The company sought to cater to those end customers looking to leverage a lower storage tier for cost savings, with a focus on power and cooling conservation. COPAN, recognizing that the data stored on a lower storage tier would be accessed rarely, developed a system architecture where only 25% of the disks would spin at any given time. Understanding that access times would be worse than traditional tier 1 storage, but significantly better than tape, the company gambled that the cost savings in both dollars per gigabyte and power and cooling costs would more than overcome the decrease in performance.

Other companies, like Data Domain, sought to boost the effective use of a low-cost storage tier by intelligently eliminating highly duplicated data sets.

Tiered Storage Today

Most other system OEMs have decided (at least for now) that end customers would much rather have the performance enabled by continuously spinning disks and have additional features integrated within their capacity-optimized arrays that broaden the appeal of two-tiered storage.

These features fall into three main categories:

- ☒ Boosting the effective utilization rate (data deduplication, space-efficient copy, wire speed compression, and thin provisioning)
- ☒ Maximizing performance through massive parallelism of spindles and increased processing power, additional cache, and intelligence data placement
- ☒ Maximizing capacity density while reducing power consumption (spin-down/spin-slower technologies, storage blades, and leveraging small form factor disk drives)

Tiered Storage Needs

Perhaps the greatest need within tiered storage architectures remains intelligent and automated data placement and movement across the various storage tiers. No IT manager desires to have transactional data sitting on a MAID array, while five-year-old emails consume its highest-performing (and usually most costly) arrays. Moreover, datacenter managers loathe the idea of managing (or more likely paying a storage administrator to manage) terabytes of data among the various tiers manually. It would be costly at best, prone to human errors, and impractical as data pools expanded into the multiple-petabyte range.

At the root of being able to manage data effectively across storage tiers, either manually or automatically, is the ability to identify the data correctly, establish and manage policies against the data, and ensure that the targeted storage tier is appropriate and available based on the business objective. If not the largest, this gap is at least one of the largest Achilles' heels to the adoption and realization of the benefits of tiered storage en masse.

Another need that hinders the adoption of tiered storage environments is the limited use of virtualization options across heterogeneous storage environments. Recent survey data suggests that end customers desire less complex and highly utilized storage environments. Some consolidation onto homogenous systems already is taking place, advancing the need for more simplistic and efficient storage virtualization and data movement capabilities. This desire is one of the main reasons why storage solutions from companies such as 3PAR, Compellent, and Dell, which embed sophisticated virtualization functions within their base systems, are sustaining strong sales growth.

FUTURE OUTLOOK

Tomorrow's Tiered Storage

The key to understanding the future of tiered storage is to recognize that storage requirements will be much more diverse. Today, most tiered storage solutions are based on systems designed primarily to support I/O-intensive and block-based applications. Most tiered storage solutions are based on the use of different disk drives and/or controller architectures to balance price versus performance.

In the future, companies will require storage environments designed to address many different information sets, ranging from large data warehouses to large pools of infrequently accessed files or records. In any given use case, the functions required in different tiers will vary radically.

At the most basic level, tiered storage solutions will vary based upon whether they support block-based or file-based environments:

- ☒ **Block-based storage:** To tier or not to tier will be a continuing refrain in the block world over the next five years. As noted earlier, intelligent movement of active data (as opposed to copies) is difficult. IDC expects to see continuous battles between solutions like Compellent's Storage Center and Dell's PS6000 that use processing and virtualization to enable tiering versus alternatives such as IBM's XIV or Sun's Amber Road that leverage memory and massive parallelism of low-cost disks to achieve higher performance and lower cost.
- ☒ **File-based storage:** Over the next five years, tiering in file-based environments will diversify dramatically and grow rapidly as file-based storage increases to almost 75% of all storage capacity shipped by 2012. Tiered storage solutions designed to support video surveillance will have a very different profile compared with a solution designed to support cloud-based backup or one designed to support the delivery of music, movies, or gaming software to consumers.

The other key developments will be more specific classification of tiers into three main types. These types will not be the traditional primary (tier 1), secondary (tier 2), and tertiary (tier 3). In many content companies, the primary storage in terms of capacity, cost, and need to manage will end up being what today is called the tertiary tier. Going forward, tiers will reflect the main function of that storage pool:

- ☒ **Performance tier:** The performance tier includes storage pools designed to deliver high levels of performance, though the definition of performance will be more variable (random I/O versus sequential, as well as primarily reads, primarily writes, or random read/writes). The performance tier will also require the highest levels of reliability, though reliability will have a different meaning in different contexts. For example, reliable storage for random read/writes in a database application is very different than the reliable streaming of a movie in an IPTV application. This is the segment where technologies such as 2.5in. SAS, SSDs, high-speed network links, and dynamic use of multiple active copies will be deployed most quickly.
- ☒ **Replication tier:** The replication tier includes storage pools designed to store copies of data/images for reuse cost effectively. The most obvious use case for this tier is to support the shift to disk-based, rapid data/application recovery. Others include rapid provisioning in test/development, data mining (e.g., Oracle's recently announced Exadata Storage Server), and ediscovery repositories. The focus of this tier will be on developing storage solutions that ingest, replicate, and store the maximum amount of replicated data at the lowest cost. Storage solutions in this tier may have very specific performance requirements (e.g., rapid ingest of backups), but most of the focus will be on adding processing performance to enable capabilities such as compression, data deduplication, or space-efficient snapshots/clones.

- ☒ **Archive tier:** The archive tier includes storage pools designed to provide *active*, though not necessarily high performance, access to large amounts of data/information for very long periods (years and/or decades). "Active" is the key to distinguishing this archive tier from the "archiving" most companies historically did with tape libraries. This tier is almost exclusively file based (e.g., database archive files, telemetry data, medical records, and personal photos). Any single file may never be accessed, but all files must be accessible in terms of being readable and findable. Not surprisingly, data density and operational efficiency (e.g., reduced power consumption) will be key functions required in these solutions. This tier will also be considered for public cloud storage services. Of equal importance, however, will be support for advanced information management capabilities, ranging from massive, clustered file systems to integrated data classification/search/analytics.

In tomorrow's tiered storage environments, the space between the tiers will be just as important as the tiers themselves. The underlying fabric/network connecting these different tiers will play a critical role in moving information between these tiers transparently (e.g., repopulating data to performance tiers in a disaster recovery scenario or shifting an old movie from the archive to performance tier due to an unexpected spike in demand). It will also play a critical role in delivering certain common services such as data encryption, data reformatting, or geographic data dispersal.

IDC's Expectations for Tiered Storage

Given the current economic conditions, it is paramount that storage administrators gain efficiencies in their approach to data and information management. Most progressive organizations recognize that there is value in the data that they preserve (almost irrespective of age). Predictive analysis on historical data to achieve models that enable growth and profitability is a worthy exercise for any organization.

However, simply keeping data for the sake of keeping data does not automatically translate into deriving value from the data. There must be context given to the data. Hence, the ability to classify data is vital.

Finally, logic suggests that a strong relationship exists between storage growth by tier and the workloads of the storage environment. Understanding these relationships will be paramount for storage OEMs and customers to rightsize their storage investments by the tiers necessary to meet the capacity and performance requirements of their organizations.

ESSENTIAL GUIDANCE

Guidance for Vendors

- ☒ Many of the IT executives currently driving tiered storage decisions are not yet fully aware of the looming archive problem, especially in the file-based storage environments (e.g., video surveillance). Sales efforts should be extended beyond traditional IT buyers to encompass information/digital asset stakeholders.

- ☒ The development of tiered storage is still in its infancy, with large growth potential and much future diversification. Use 2009 as a year to align your product portfolios and features with high-priority customer needs and desires via aggressive development, strategic partnerships, or acquisitions. This effort should include consideration of cloud-based service offering and/or partnerships.
- ☒ No one tiered storage solution will address all use cases. Prioritize near-term sales efforts based upon which opportunities best fit with current products/sales channels. Identify larger long-term opportunities and start investing in technologies and partnerships that will address the bigger needs more effectively.
- ☒ Customers are looking for simplicity and efficiency, as well as help with justification. Migrating to new platforms can be disruptive and, worse, disastrous, resulting in dreaded downtime. Ensure that you develop migration tools, as well as future migration paths, within your own product portfolio to ease integration and growth.
- ☒ Develop realistic tools to help customers understand the likely return on investment — and then dynamic tools to help track the efficiency of their storage environments. Assist customers in identifying and developing key performance indicators that will allow them to equate return on technology investments with business objectives.

Guidance for Customers

- ☒ Take time to learn the various technologies available to you and how the technologies apply to your specific organization.
- ☒ Make a concerted effort to find the "undiscovered" pockets of data/content that are quickly, but stealthily, becoming major storage consumers. Areas to watch include video recording/surveillance, ediscovery, and any business activity or application that is converting to a digital format. Assess current and future data growth rates as well as long-term archiving needs for these fast-growing content repositories, and start thinking about how well these mesh with current datacenter plans.
- ☒ Develop a well-thought-out holistic approach that addresses both near-term and long-term storage requirements, rather than a point solution that addresses a specific need that may or may not evolve well with your storage environment.

LEARN MORE

Related Research

- ☒ *Achieving Storage Efficiency Through Storage Tiers* (IDC #217734, March 2009)
- ☒ *A Compelling Use Case for Solid State Disks* (IDC #cUS21763009, March 2009)

- ☒ *The Economic Meltdown and Its Chilling Effect on the Cost to Power, Cool, and Manage Enterprise Storage* (IDC #217158, March 2009)
- ☒ *Make Room, Server Virtualization: Collaboration Becomes a Major Storage Challenge in the Enterprise* (IDC #lcUS21657209, February 2009)
- ☒ *EMC's New Celerra Systems — Improving Storage Efficiencies* (IDC #217070, February 2009)
- ☒ *Worldwide Storage 2009 Top 10 Predictions: Grappling with Content Growth in a Contracting Economy* (IDC #216026, January 2009)
- ☒ *Storage in the Cloud: What, How, and Who?* (IDC #215052, November 2008)
- ☒ *EMC's Atmos: Making Rain in the Cloud* (IDC #215169, November 2008)

Synopsis

This IDC study evaluates today's tiered storage environment. What is the state of tiered storage? How is this technology evolving? What do the solutions look like today? And how are future requirements shaping the tiered storage solutions of tomorrow? This document answers these questions and considers whether or not the industry has "tiers of joy" or "tears of despair."

"Today's economics, a growing diversity of data types, and the desire to manage an increasing amount of data more efficiently are driving IT managers to evaluate more sophisticated and efficient storage solution technologies. A key technology that is destined to assist in this endeavor is tiered storage." — David Reinsel, group vice president, Storage and Semiconductor research

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights.

Copyright 2009 IDC. Reproduction is forbidden unless authorized. All rights reserved.