



## Why Storage Tiering Still Matters. A Lot.

Storage tiering helps enterprises deal with data growth more cost effectively.



By Christine Taylor

Tiered storage is the process of assigning progressively less expensive storage categories to progressively less valuable data. It's up to IT to classify storage tiers using a matrix of performance, price (Capex and Opex), storage capacity, and data services. Classifying data priority is not entirely up to IT. Within the same storage system, automated tiering functions will classify data by features like I/O patterns and move it accordingly within the storage system's internal storage tiers.

However, IT will need to assign data priority by business need in order to effectively migrate data throughout the storage infrastructure, ultimately landing in highly cost-effective cold storage. Different companies will assign different data priorities according to their business and compliance needs. For example, aging is the most common metadata for demoting data to less expensive storage, but other factors may affect the outcome. For example, IT may progressively demote aging data and eventually add it to cold storage on tape or cloud. But some aging data may reside long-term on SATA on-premises storage because it is subject to regular information audits.

An effective overall tiering process largely depends on automating data movement across the storage infrastructure. If IT is spending a lot of time manually tiering data — or not tiering at all because it takes too much time— then the company will not see value out of storage tiering. Automated storage tiering across data lifecycles is a real necessity for getting the kind of value that IT needs out of long-term data storage.

### **The Difference between Tiering and Flash Caching**

Let's talk about a perennial point of confusion: the difference between tiering and caching. When you're talking about tiering data through the entire storage infrastructure, there is no question that you're talking about storage tiering. But when you're talking about dynamic data

movement on a single storage system, then the difference between tiering and caching is not as obvious.

The big difference between the two is that caching copies or mirrors frequently accessed data to flash drives, whereas tiering moves data. There are of course distinctions between caching products: some are block and some are file-based; some operate on the server side and some from the storage system; some have only a single flash tier and others contain internal SSDs for storing cached data. But in general, cache products copy data to a high performance flash tier to accelerate applications.

Dynamic storage tiering within the array does not copy data, but moves it between storage tiers based on access patterns. If all you are after is I/O performance in a single system, then the caching model may work well for you (or consider paying for an AFA). However, if you want both accelerate performance and realize more cost-effective storage usage, tiering will go farther and accomplish more. Unlike a caching layer that keeps flash media otherwise unusable, tiering will deliver data to faster flash tiers for high performance processing, but unlike a cache tier the storage is usable and not just a copy of the data. Then as data access patterns change, the tiering software will move it to less expensive storage tiers for cheaper retention.

## **Storage Tiering Domains**

Vendors offer automated tiering in very different locations and configurations. The most common choices include within the hypervisor, within a single storage system or connected same-vendor systems, or across the entire storage infrastructure. Most companies will invest in all three approaches depending on their needs.

### ***In the Hypervisor***

Hypervisor-based storage tiering optimizes VM performance by dynamically placing them in differing storage classes. High-performance VMs reside on high performance storage tiers, while VMs that don't require lowest latency reside on less expensive tiers.

[VMware](#) Storage Dynamic Resource Scheduler (Storage DRS) typically manages live tiering across the VMware environment. DRS recommends placement based on I/O loads and storage constraints based on reducing latency impact. Default placement is every eight hours but admins can customize the schedule and assign storage classes.

[Microsoft](#) uses features in Windows Server 2012 R2 and Hyper-V Server 2012 R2 to automated storage tiering. This process continuously analyzes I/O and moves blocks accordingly between flash and hard disk tiers.

### ***Within a Storage System***

Storage tiering does not only cover moving data from one storage system to another. Storage tiering may also refer to dynamic data movement within a single storage system based on application performance needs. In this case, storage tiering is primarily concerned with accelerating data performance within the production storage system, and secondarily storing aging production data on less expensive tiers within the array, or storing snapshots on cheap storage.

This technology tiers data within a multi-tiered single storage system, or between connected single-vendor systems. The most common configuration is a hybrid storage system where the automated tiering function stores high priority workloads on a flash layer, then moves them to hard disk tiers. In reality this process may also occur on all-flash arrays or on hybrid arrays where there are two classes of flash that exist in front of the hard drives.

Dynamic automated tiering within the system is a hallmark of [Dell Storage Center](#). SC Series Volume Advisor proactively monitors data placement and storage optimization throughout all connected SC systems. Flash systems with high performance tier 0 and automated tiering services include [IBM hybrid Storewize](#) systems with Easy Tier. Easy Tier analyzes real-time usage analysis patterns and assigns data to different Storewize tiers.

### ***Storage Infrastructure***

This is the process of tiering storage across different classes of storage systems. In the recent past it was common to have a 3, 4, or even 5-stage tiering infrastructure consisting of a high-performance array, a nearline disk drive, and tape. This is still a common solution because it works and because there is a large existing investment.

However, given large rates of data growth and rising storage expenditures, tiering choices across the infrastructure have become far more complicated than they used to be. There is still a minimal 2-stack core in every infrastructure tiering configuration, which is primary storage to tape or cloud. A 3-stack core is probably more common but may not be the traditional one of primary storage system, nearline disk, and tape.

Today the 3-stack core may consist of a flash tier and a hard drive tier on a single production system, an active archiving tier on disk or high performance tape library, and a long-term data retention tier on tape or in the cloud. Even within this updated configuration there will be alternate architectures. For example, your top production tier may be an AFA with a Tier 0 high-performance flash layer, and software that dynamically tiers data between Tier 0 and Tier 1 SSD drives. Or you might have a hybrid flash system with Tier 0 and Tier 1 flash layers, and internal hard drives classified as nearline Tier 2.

There are additional options for a nearline Tier 2. A high-performance tape library may act as a nearline storage system for active archiving. So might a cloud-based tier that is engineered as a Tier 2 with fast recovery options and additional data services. [Microsoft Azure StorSimple](#) for integrates on-premise appliances with a fully active cloud tier containing data services and failover functionality.

Both tape and cloud may serve as Tier 3 or 4 cool or cold storage. Major public cloud offerings for long-term data retention include cool storage service [Google Nearline](#) for long-term storage with reasonable recovery times, and cold storage [Amazon Glacier](#), which has recovery options but is geared towards retaining data with no need of regular recovery.