

STORAGE SWITZERLAND REPORT

Maximizing Tiered Storage ROI: Performance-based Storage Tiering for Fibre Channel SANs



Eric Slack, Senior Analyst

Costs of high performance, fibre channel SAN infrastructures can be significant and storage tiering is often employed as a strategy to control those costs. However, most tiering methodologies simply prioritize groups of data based on access patterns and assume that the highest tier data should be stored on the fastest storage arrays. This assumption is inaccurate and typically results in significant under-utilization of the most expensive storage.

Performance-based storage tiering seeks to address this problem by measuring the performance actually needed and adjusting the choice of storage hardware to the applications for each tier's data, maximizing the ROI of tiered storage systems.

Storage tiering is an optimization strategy that strives to match data sets with storage systems, in an effort to reduce overall costs. The concept holds that some data is 'over-provisioned', or likely residing on higher performing storage than is necessary and increasing the utilization of storage is the way to reduce costs. To the extent that some of these data can be moved off to lower-performing storage, the costs/GB of

capacity for the system as a whole will go down - significantly.

Storage tiers have been implemented with different drive technologies; fibre or SAS drives being the upper tiers and SATA drives being the lower tiers. Recently SSD technology has added a 'tier 0' above the original tier 1 but the concept is the same, except that the migration is upward to the higher performing tier. Classifying data for each tier has been less precise, but has generally been driven by how often a data set is accessed.

Storage tiering assumptions

In most tiering environments some cost savings are realized, just by moving obviously inactive or lesser important data off of tier 1 arrays. Moving more active, more important data, on the other hand, can be more risky since performance is often critical for the applications associated with these data sets.

Since organizations don't usually have detailed information about the performance these applications really need, most assume their tier 1 data needs to be stored on their fastest, tier 1 storage. The result is an over-provisioning of storage performance to these data sets and an under-utilization of those resources. This, in turn, drives up storage costs including the associated environmental costs; power, cooling, floor space. As an example, if you could move the data from 300GB 15K RPM FC drives, typically deployed in a mirrored RAID configuration, to 1TB 7200 RPM SATA drives in a RAID 5 configuration, overall storage costs, including environmental would decrease by about 80%.

Current infrastructure monitoring systems focus on providing status data about components in the SAN - storage subsystems, network switches and HBAs. This information can identify which data sets are the most frequently accessed and which applications they involve, and can indicate which data probably needs the most performance. But they don't provide any information on how fast data needs to be or can be supplied to servers. At best, they enable data sets to be prioritized, telling IT which data needs to be the most readily available, but not specifically how available. Without this information, the performance needed by the storage arrays that comprise a tiered storage system is left to assumption. For tier 1 data sets, given the importance of the applications they serve, the fastest, most expensive storage arrays are typically used, storage which is often many times faster than necessary. In more rare cases this top tier storage could be too slow, indicating a need for SSD.

The point is that current tier 1 storage is often not the optimum placement for tier 1 data.

While tier 1 data may need much faster delivery than tier 2 data, the reality may be that current tier 2 storage is fast enough for that data. But without accurate, real-time information on just how fast is 'fast enough', storage managers often assume they need tier 1 storage performance.

Don't rely on assumptions

Performance-based storage tiering uses real-time measurements of storage network data transactions to determine the performance needed for application servers in each tier. Then, storage types are chosen that support those data rates, ensuring that storage resources aren't under-utilized. This concept challenges the assumption that an organization's tier 1 data automatically needs the performance of the fastest storage available.

Performance tiering measures real-time application latency, or the amount of time applications wait for the storage system to provide data, through network-connected, physical-layer traffic access points, or TAPs. These passive, out-of-band appliances capture network traffic without affecting performance. Systems like VirtualWisdom from [Virtual Instruments](#) take this storage transaction data and display the activity between servers and storage, including response times, queue depth and average utilization rates for each storage resource.

Through the collection of these transaction times and some analysis, a baseline can be established, revealing how fast data is being delivered to the most critical applications. Then, it can be determined if the storage subsystems are in fact slowing applications down, or if the opposite is true, that they're running below rated speed due to bottlenecks elsewhere in the storage infrastructure. Storage tiers can then be created on platforms that provide the appropriate levels of performance for these data sets, based upon accurate, measured information about the needs of the applications, instead of assumptions.

The data provided can also help identify bottlenecks in the SAN infrastructure, not just the storage, that are slowing transaction times down. In addition, this SAN monitoring can be used to test modifications before they're rolled out to production servers, ensuring that steps taken will be effective. After storage response time is optimized for each application, trending and more detailed analysis can help with infrastructure design and upgrades.

Improved storage tiering

Traditional storage tiering classifies data and places it on the most cost-effective tiers of storage. Performance tiering doesn't replace but augments this process by providing the information required to adjust the performance levels of storage in each tier to improve utilization. Performance tiering should be a prerequisite to storage tiering, as it provides the data needed to provision storage tiers accurately, reducing the waste of under-utilization. It also enables storage managers to assign new data sets to an appropriate tier with confidence, even for an organization's most critical applications.

Storage tiering is an important strategy for optimizing storage expenditures. While the concept is sound, tiering methods traditionally store the data supporting an organization's most important applications on its highest performing storage systems. Many SAN environments can't supply data to these applications as fast as their top tier storage systems can provide it, and this 'tier 1 storage for tier 1 data' assumption results in an under-utilization of these expensive storage resources.

Performance-based storage tiering uses SAN-connected access points to measure real-time 'data transactions' and determine the storage performance needed by application servers. This information is used to construct storage tiers with an appropriate level of performance, resulting in greater storage utilization and a maximum ROI for tiered storage infrastructures.

About Storage Switzerland

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>.