

STORAGE SWITZERLAND REPORT

SSD OR AUTOMATED TIERING?



George Crump, Senior Analyst

Automated tiering solutions decide, based on access patterns, if a set of data should be on Solid State, Fibre or SATA based storage. The value to storage administrators is they can take the supposed work out of implementing multiple tiers of storage. As a result traditional storage suppliers and a few start-ups are adding automated tiering capabilities to their storage systems or coming out with new systems to support this capability. These automated systems do add a another component to purchase and manage in the environment. This leaves IT managers with a decision to make when faced with performance problems: implement a separate solid state disk (SSD) tier or implement automated tiering.

The first step in making the decision is to understand the nature and scope of the performance challenges. Automated tiering is a broad-brush approach attempting to provide a performance boost across the enterprise; which, as part of that process, will likely leverage SSDs. A specific SSD drive, PCIe SSD card or SSD system is more surgical in nature, is designed to address a specific performance issue. Storage I/O performance problems, for many organizations, boil down to just a few applications or data sets. For these organizations, it may be less expensive to implement SSD and forgo automated tiering altogether.

One of the goals of automated tiering is to try to improve performance or present adequate performance at a lower cost. To do this most automated tiering systems will also

leverage SSD. That requires an understanding of the impact of SSD on the environment and what SSD requires of the environment, to deliver on its performance potential. As we describe in Visualizing SSD Readiness this means that the application must be able to generate enough simultaneous I/O requests to warrant the zero-latency response of an SSD system. In addition, to take full advantage of SSD also requires that the entire I/O chain be able to deliver data as fast as the SSD can. This is the number one issue with an effective automating tiering strategy.

SSDs are universally recognized as dramatically faster than disks, with an SSD capable of delivering the I/O per second performance of hundreds to thousands of HDDs. When many SSDs are put behind a storage controller or appliance, the amount of work that appliance has to do to be effective is often dramatically underestimated. At a minimum, for each I/O a host I/O and a storage I/O must be generated. This means the tiering controller has to handle double the amount of IOPS that either the storage or the servers are handling. Since automatic data migration is part of this equation as well, there must also be additional I/O to move the data to the tier and off of it as well handling any other storage management features (RAID, shapshots, thin provisioning, etc). This can make the storage controller the performance bottleneck, and understanding how the storage controller can scale to handle I/O loads is the question that needs to be answered for a successful automated storage tiering strategy.

While many enterprise environments have a few servers that can generate the level of I/O requests that make a storage controller the bottleneck, the whole value proposition behind automated tiering requires centralized storage leveraging a unified pool of storage tiers. This easily makes the storage controllers the bottleneck and dramatically limits the performance scalability.

In addition to having the appropriate storage controllers to benefit from automated tiering, it's also important to understand that the technology requires that something "new" be implemented in the enterprise. This may either be in the form of a storage hardware upgrade, a storage software upgrade or a standalone appliance. Some automated tiering solutions require changes in how the servers connect to storage, so there may be some additional implementation challenges as well. In any case, there is likely to be some education required and some storage tiering policy decisions to be made as well. Some storage manufacturers will also charge for the automated tiering software, which, as a license, will probably need to be updated annually.

The investment in automated tiering may make sense if the organization can see a need for a wider access to SSD or more likely is trying to reduce storage costs by tiering to cheaper media, where the goal isn't to increase performance, but to maintain it at a lower cost. This is especially true if it's time for a storage refresh. If neither case is true then pinpointing a performance problem with a single SSD appliance may be more cost-effective, and will deliver a much higher performance gain to the specific critical application.

With automated tiering solutions provided by the primary storage supplier, the storage hardware for all the tiers is typically sourced from that vendor. The customer is locked into that vendor more so than ever. This has ramifications not only on the high performance tier but also on the high value tier. The primary storage provider may not have chosen the best type of storage for all environments. For example, in the high performance tier it may be a write-

heavy environment that's causing the performance problem. And the solution may be more suited to DRAM based or Flash SSD with advanced controllers, which the primary storage vendors are not using.

Primary storage suppliers use flash technology that's in a drive enclosure form factor and they typically need a minimum of four to six of these drives for optimal performance and reliability. This means minimum capacity configurations of almost 1TB of SSD. For many environments that's far too much capacity at the SSD tier and that, despite the price drops in SSD, the technology is still at a 10 to 15X premium of traditional mechanical hard drives. For many environments a few hundred GBs of SSD is all that's required. Additionally, the components that surround the flash memory to make it drive-plug compatible add to the overall cost. SSD systems are purpose-built to house SSD and don't require the more expensive design of modules made to house disk drives. Also, as mentioned earlier, they don't need a minimum amount of capacity for optimum performance or redundancy. As a result, not only are SSD systems typically less expensive than comparable automated tiering systems, they often require only half as much capacity.

Another issue with automated tiering may be the automation itself. Many IT professionals know exactly what data and servers are causing I/O contention issues and the business units know which application's performance matters. Frankly, they just don't need the automation. Again, implementing SSD to a specific application may be a quicker, cheaper and simpler route to solving the performance problem. While this would appear to require a new implementation, it's only for a specific data set. As we discuss in "[Implementing SSD and Maintaining Disaster Recovery](#)", standalone SSD can be implemented without requiring a re-architecting of the backup process. The cost of a single SSD system may be less expensive and less disruptive than buying an automated tiering-enabled storage system, especially if that involves buying an entirely new storage system.

There are also performance advantages to a purpose-built approach to the solution. While a few automated tiering systems are separate, dedicated appliances, the solutions provided by the primary vendors mix this technology in with all the other functions that the storage controller is responsible for. In these cases, the storage system has a very mixed workload that it has to deal with, including services like provisioning and data protection. A separate SSD tier means that high I/O traffic can now be offloaded from the primary storage controller, freeing it up to provide better performance to the other servers in the environment as well as data services. The result is a distributed I/O workload instead of forcing all the I/O through a single controller group.

Finally, a purpose built SSD system can still be integrated into an automated tiering infrastructure at a later date, if the capacity is justified. The storage manager can leverage block virtualization appliances and/or file virtualization appliances to use all or part of the SSD appliance in an automation strategy.

Automated tiering is a popular and important technology to fully leverage all of the storage tiers available in an environment. When the timing is right many data center managers may choose to implement this capability. It is important to understand though, that automated tiering is not required to take advantage of SSD. If you know what specific data needs can take advantage of the SSD performance boost, a purpose built SSD solution may be more cost effective and a better short term strategy, especially if that data represents a relatively small percentage of the enterprise's overall data set.

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>.