

STORAGE SWITZERLAND REPORT

START WITH AUTOMATED STORAGE TIERING TO LOWER NAS COSTS



George Crump, Senior Analyst

For some storage architects, automated tiering is thought of as the last step in NAS performance tuning and is not a consideration when designing or deploying a new NAS infrastructure. In reality, they should be looking at automated storage tiering upfront to improve performance and lower overall NAS costs rather than after new gear has been purchased and deployed. Doing so may allow them to purchase a lower cost, mid-range NAS system instead of investing in a more expensive, higher-end NAS system. But first, let's take a look at the historical choices and trade-offs the architect has had in designing a next generation infrastructure before we examine the benefits of the newest approach via automated storage tiering.

Performance in NAS systems has come from two areas: the amount of hardware resources that are used and the quality and capability of the selected software. Obviously the higher the investment is in either, the more expensive the overall solution becomes. From a hardware perspective, performance has typically come from investing in more powerful processors, more and higher bandwidth network interface cards and more and faster disk drives. But this approach can introduce greater risk into the

environment. Larger drive counts in RAID groups does improve performance but can increase the likelihood of multiple drive failures, causing unrecoverable data loss. Snapshots of data volumes may increase data recoverability, although not from a failed RAID set, but in many systems can incrementally reduce overall system performance.

From a software perspective, the most cost effective file sharing solution is to leverage standard operating systems, especially open source software like Linux. These operating systems can typically present both NFS and CIFS file shares, but sometimes does so at a reduction in 'usability', fewer software features and less scalability, when compared with a purpose-built NAS OS. Purpose built NAS operating systems like ZFS, FreeNAS or NexentaStor offer greater functionality for their relative higher cost, but still run on off-the-shelf server hardware. The highest end solutions marry the hardware and software components together for an integrated, but very expensive, solution. Conversely, a software-only NAS operating system is less expensive, but may not be as efficient and may require a greater investment in hardware.

The ability to get the best attributes of all these solutions - high performance, full data services and reasonable cost - may be had by considering an automated tiering device early in the design of a new NAS system. Appliance-based automated tiering systems (ATS) are independent from the storage itself and can work with a mixture of vendors' hardware and NAS OSs. When designed in, as part of a new NAS strategy, an ATS provides the organization with tremendous flexibility in the NAS selection process. With this configuration, most active data is moved from the NAS system to the ATS. When that data is accessed again it's served from the ATS which is typically a higher throughput device, like RAM or Flash, than the NAS head.

When trying to drive down the cost of the new NAS systems yet maintain performance, pay special attention to 'persistent' ATSs like those offered by [Avere Systems](#). These devices are much more than a traditional read or write cache, which effectively buffer data for these I/O processes to increase performance. A persistent ATS stores data within the system and, while it does act as a cache for both reads and writes, it won't flush until the new data is needed on the NAS system. This flush to NAS can also be scheduled to coincide with snapshots and backup windows. The persistent ATS systems, unlike other forms of caching, typically have a mixture of storage media themselves including RAM, SSD and high speed SAS drives. This provides for a complete internal tiering of active data as well as a migration of inactive data to the NAS.

In the persistent ATS model, the role of the NAS filer changes from being responsible for serving up data as fast as possible to managing and protecting that data. This allows the customer to spend less time and money tuning for optimal performance and become more focused on data protection and management. The selection of the NAS backend now can focus on what the customer needs in these areas. If those needs are somewhat basic, for example, possibly an off-the-shelf Linux fileserver will

suffice. If so, then the cost savings on the NAS backend to this environment will be enormous. If greater data services are required then an investment in a NAS software option that can provide some of the needed data services yet still run on industry standard hardware may be the best use of budget dollars.

In either case the investment in the NAS hardware can be greatly reduced. First there is no need for extensive network connectivity, just enough bandwidth to feed the persistent ATS and be able to receive that data when the cache flushes. Many customers find that they don't even need 10GbE cards in the NAS repository as the high performance I/O needs are consolidated at the persistent ATS instead. Additionally, the investment in the storage system is greatly reduced. Since the NAS is now in the role of 'data protector', high capacity drives, even SATA drives can be used. This not only lowers the investment in storage hardware but also the costs associated with powering that hardware.

Good evidence of how well a persistent ATS can drive performance up on midrange systems as seen in Avere Systems recent performance on the SPECsfs2008 Industry benchmark. A simple two-node Avere cluster achieved an Operation Rate of 43,796 with only 1.33 latency, using only one file system and 26 disks. The NAS backend for this test? - a simple Linux server. This result placed it comfortably at the top of the higher end NAS systems. Scaling the Avere system to 6-nodes, still with a single file system, raised the operation rate to 131,591 with only 1.38msec latency, putting it at the front of the pack.

While real-world performance mileage will always vary, it's almost assured that if there is a performance bottleneck in the environment, persistent ATS should be able to help. In a new system purchase the ATS can enable the use of significantly less expensive NAS hardware and software driving cost out of the initial and subsequent NAS purchases.

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>.