

STORAGE SWITZERLAND

Storage-at-Scale Systems Save Operational Costs



Eric Slack, Senior Analyst

Files comprise the bulk of data growth and file systems are generally considered easier to use than block storage. These two factors are pushing the adoption of NAS storage and driving the need for larger and larger file storage infrastructures. This is especially true in high performance applications, such as media and entertainment, life sciences, large scale virtualized environments, energy, etc. compared with legacy file sharing use cases. But some NAS architectures can suffer an increase in cost per TB as they scale. Storage-at-scale systems can save operational costs by maintaining storage density and controller efficiency as they expand, while keeping administrative overhead in check.

Storage Density

Storage density can be expressed as the amount of physical resources (space) required to produce a given capacity of storage - basically how much usable capacity per rack unit (U) or floor tile in the data center. The caveat here is that this capacity must provide a consistent, acceptable level of performance as it grows. A traditional disk array with a single or dual controller architecture and

disk drive shelves, for example, can scale linearly by adding drive shelves, but most come with a fixed amount of processing power in each controller. While fine for many use cases, this architecture can cause performance problems (especially IOPS) for some workloads, as spindle counts grow. This means having to buy additional physical storage systems before the maximum capacity of existing systems is reached, resulting in lower storage density.

Scale-out storage systems were designed to address this issue by combining processing power and network connectivity into storage modules or nodes. This corrected the performance problem that traditional disk arrays faced as they scaled, but caused another problem with density. As these clusters grow to provide storage capacity so does their node count, consuming more data center space with additional chassis. Traditional NAS systems typically don't fare any better. With an fixed controller and disk shelf architecture, these systems often suffer from the same drive density problems that traditional disk arrays do. And, adding more discrete systems to the environment to keep up with capacity demands can create a NAS 'sprawl' issue.

A storage-as-scale system, like [BlueArc's Titan and Mercury](#), can scale processing independently from disk space, allowing the system to add controller horsepower as needed to support more spindles. This 'scale up' architecture can supply the performance needed for each level of capacity, maintaining storage density as it grows. Compared with multiple NAS systems or a scale-out storage environment, this reduction in hardware (fewer chassis, controllers and connectivity) means less spent on power, cooling and data center real estate. It can also mean fewer storage 'silos', reducing ongoing costs of maintenance, licensing and administration time to support these systems.

Controller Efficiency

Hardware and software design also drives data efficiency which impacts storage density. A system that can produce more IOPS and throughput for a given level of CPU output can support more disk drive capacity with each controller, which means fewer controllers or controller upgrades are needed for a given storage capacity and less physical space and power are consumed. One factor that affects controller efficiency is the use of discrete hardware, like ASICs and FPGAs, for storage and file system operations. Systems that perform protocol processing (CIFS, NFS, iSCSI, NDMP) in software consume system CPU cycles, cycles which aren't available to support storage capacity. Using dedicated hardware for these overhead processes allows an at-scale system to support more spindles per storage controller, increasing density.

On the software side, controller efficiency is affected by how a system handles metadata and file system overhead. Metadata (file attributes, permissions, access histories, file system indexes, etc.) and "metadata operations" are present in every storage I/O process, but also in functions like replication, deduplication, snapshots, etc. In response to this, some storage-at-scale systems are designed to

handle metadata more efficiently, like creating a two-tiered architecture and pinning active file metadata in cache, or on the fastest available storage. As explained in this Storage Switzerland report, [tiering metadata](#) storage provides more IOPS for processing metadata and keeps slower disk storage from becoming a bottleneck for system performance.

An object-based file systems (OBFS) can also improve efficiency and flexibility, particularly in systems that need to scale to very large capacities. The object architecture organizes data as a flat collection of unique object ID numbers, not a complex hierarchy of files, folders and directories (metadata) like traditional NAS architectures use. Without all this metadata, [OBFSs can also support much faster replication](#) than can either file or block-based methodologies, which can save significant amounts of time for storage administrators in the data protection process. They're also much more efficient, resulting in higher performance, less CPU and less storage overhead, especially as the system grows.

In addition to raw performance, a storage system should scale storage services, like snapshots, replication, cloning, etc, to keep pace with capacity growth. If snapshot space runs out before physical capacity, the effective storage limit for the system is reduced. If available processing power can't keep up with the demands of off-site replication as the system grows, the same thing occurs. The bottom line is that storage services might be as important as read and write performance. Maintaining these functions can be a determining factor in how large a storage system can effectively scale. Storage-at-scale systems that optimize metadata operations can support these storage services as capacities grow, improving density and lowering operational cost.

Administrative Overhead

Administration is another factor that impacts storage costs. The time required to conduct regular 'care and feeding' activities such as patching, updates and troubleshooting, and the effort required to expand a storage system are both determinants of its operating cost. Generally, a scale-up system that can expand by simply adding disk drives will consume less admin time, compared with adding storage nodes to a scale-out system. This is even more true with adding NAS systems which often includes data migration to balance workloads across multiple discrete storage units.

At the software level, storage-at-scale systems leverage a number of technologies to reduce administrative overhead. Virtualization and thin provisioning can be used to create storage pools from multiple external arrays or multiple internal disk tiers and dynamically allocate storage to improve capacity utilization and performance. They can also simplify end user storage management by creating virtual volumes for different business units or applications, each with its own cluster namespace and access controls.

With the growth in server virtualization, especially using NFS, storage at scale systems should also support integration with VMware's vCenter. Using the vCenter APIs, a storage at scale system, like BlueArc, can dramatically reduce administrative overhead. Tasks like managing snapshots to support the backup and recovery process of VMs can be accomplished from within vCenter.

Storage-at-scale systems can provide better storage density and maintain that density, even as they expand to very large proportions. This factor, combined with their data processing efficiency enables these systems to better support high performance, high capacity file services environments, compared with scale-out and traditional NAS architectures. The result is reduced operational costs, a characteristic that 'keeps on giving', as lower, power consumption, physical space and administration time pay dividends every month.

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>

Copyright © 2011 Storage Switzerland, Inc. - All rights reserved