

STORAGE SWITZERLAND

UNSTRUCTURED DATA GROWTH REQUIRES ENTERPRISE 'STORAGE AT SCALE'



Eric Slack, Senior Analyst

For many IT organizations data growth is continuous and substantial, especially file-based data or unstructured content. IT professionals must plan for this growth and control its impact on their infrastructures, which means managing storage growth to support data growth. To do this, file storage systems must be able to expand without affecting operational characteristics or the end user's experience. These systems must provide the same, consistent technology features and business benefits 'at scale', in other words, at the performance, capacity and operational level a user needs to support their business.

Building scalable shared storage is a simple concept that has been complicated by many storage vendors, often to compensate for the shortcomings of their systems. Designing a storage system to expand to very large proportions is possible today with several types of storage architectures. What's more challenging is designing a storage system to maintain its functionality as it grows. Many manufacturers run into problems when they're forced to expand beyond their ideal capacity range - the point at which each additional drive degrades performance. This leads to a decay in overall efficiency, resulting in a number of issues like: a degradation of performance, a surge in resource consumption, an increase in operational complexity - or all of the above.

The symptoms of scaling an inefficient storage system include having to keep buying 'silos' of NAS storage to

provide needed capacity, building clusters of scale-out storage nodes larger than needed in order to get enough processing power, expanding data center floor space and power distribution to accommodate racks of low density storage arrays or using complex storage management systems to adequately provision a diverse infrastructure. In short, scaling inefficient storage systems means more cost throughout the environment.

Storage at Scale

In order to support the infrastructure required to accommodate projected data growth, a primary storage system serving unstructured data should maintain its functionality as it scales. This is the objective that drives manufacturers to develop new and innovative systems. As will be discussed below, to achieve this objective 'storage at scale' systems should:

- Be able to increase capacity easily and scale to levels that can accommodate projected data growth
- Maintain performance (IOPS and throughput) - at scale
- Keep operational requirements (administrative overhead) within reasonable levels - at scale
- Keep resource consumption within reasonable levels (space, power, CPU) - at scale
- Keep cost within reasonable levels - at scale
- Maintain data protection - at scale

Capacity at scale

The storage-at-scale system should first be able to expand in capacity large enough to support expected data growth, and do so independent of performance. This enables the storage system to expand to the size required, but also to the 'shape' needed to support the environment. In some cases, like an archive, this shape would be very high capacity and modest performance. In others, high performance with only moderate capacity would be needed such as rich media files being edited by video game and TV commercial developers. In still other environments a linear growth curve, with capacity and IOPS/throughput increasing at similar rates, could be best. Examples of this use case could be business critical applications in virtualized data centers with 100s of virtual machines or 1000s of virtual desktops or eDiscovery projects undertaken by enterprises to support legal requirements.

In order to meet this capacity challenge, modern storage systems could benefit by having the ability to create application-specific or user-specific storage tiers to accommodate different stages in the lifecycle of a dataset. While tiered storage isn't required to support every application, for environments heavy with reference data, like the broadcast or medical imaging industries, it can make sense. In these scenarios, the system should be able to move data between these tiers transparently to users and applications on the front end. This enhanced data mobility would also enable the storage-at-scale system to incorporate third party NFS storage devices into a common infrastructure, when appropriate, to help meet capacity demands by repurposing existing assets.

In addition to volume capacity, the system should also support file system growth as well, which requires efficient metadata handling to manage the increase in data objects. For instance, search and indexing for eDiscovery is usually 50% metadata operations. Efficient metadata optimization increases the performance of file access. This efficiency with data objects and metadata also supports the increased replication and cloning activity that typically grows with storage capacity.

Performance at scale

Perhaps the most important aspect of storage-at-scale is the ability to maintain throughput and IOPS at any capacity point along the system's growth curve. This requires a greater efficiency in storage processing systems to provide more net available data processing horsepower at any given level of resource use. This kind of CPU efficiency is needed in order to support higher capacity points at a consistent performance level. Examples would be large scale email applications or use cases that need more performance overall, such as large virtualization implementations, OLTP, web 2.0 and HPC (High Performance Computing) workloads.

This efficiency can be made possible by moving the most critical, in-data path operations from file system software running on general purpose CPUs to programmable silicon hardware, like FPGAs. Parallel or concurrent 'pipelined' data processing with independent hardware stacks can eliminate data bus bottlenecks created by sharing memory or I/O bandwidth and dynamic read caching can keep files on the fastest storage as they become 'hotter'. Similarly, on the software side, using [object-based architectures](#) instead of traditional file systems, can provide the flexibility to perform these metadata operations independent from data operations. This enables the system to optimize these critical operations by pinning metadata to the highest tiers of storage and offloading its processing to a dedicated module.

Performance efficiency must also be maintained as the system scales. Some practices designed to improve performance (like writing only on selected portions of the disk) can worsen efficiency as utilization and storage density drop. Instead, leveraging storage virtualization allows new capacity (and spindles) to be added to existing storage pools, balancing read and write performance as capacity grows. File system growth must also be accommodated. An object-based file system can do this by organizing data as a flat collection of object ID numbers, instead of a hierarchy, like traditional file systems.

Operational efficiency at scale

Another aspect of efficiency at scale is operations. Storage consumes finite resources like power, cooling, data center floor space and the time of qualified storage and application administrators. Storage efficiency enables these resources to be optimized, something that must be maintained as capacity grows. Storage density and hardware utilization are important for physical resources, but administrative overhead should also be kept under control as the system expands. Administrators supporting IT-at-scale should be able to operate, maintain and upgrade the system effectively with consistent levels of input.

By maintaining high density, [a scale-up storage system](#), like [BlueArc's Mercury or Titan](#), can keep its node count to a minimum, improving data center space and energy efficiency. Some scale out storage architectures, on the other hand, may require more independent nodes, each with its own chassis, connections, power and cooling requirements. In these cases as system nodes increase, CPU or disk utilization levels typically decay as well, further reducing resource efficiency. In a similar fashion, the collection of separate NAS or storage arrays that many IT organizations accumulate can have much lower storage 'capacity per box' ratios, producing the same problems with resource efficiency.

In order to keep the amount of administrative overhead per TB under control for IT operations-at-scale, a storage system must be easy to use and require fewer steps to operate and less manual intervention when changes need to be made. Day to day management can be simplified by using a NAS system rather than block-based storage on a SAN because NAS systems manage the file system and underlying storage together. When expanding file system capacity, by comparison, a block system requires LUN operations at the array level and a separate set of file system tasks at the file level.

consolidating the infrastructure from a collection of independent NAS devices or storage arrays onto a single NAS platform, which operations-at-scale makes possible, can better support dynamic environments, like server virtualization. A larger, storage-at-scale system can require less movement of data around to balance workloads or meet peak conditions. Also, a system with the ability to increase snapshot and clone capacities to keep pace with storage growth will also help reduce overhead.

Server virtualization is a fact of life for enterprise IT organizations. To accommodate this infrastructure, enterprise storage should be integrated with virtualization platforms, like BlueArc's JetCenter plug-in for VMware vCenter. This integration can provide a single point of management for the storage provisioning, snapshot and data protection processes. Writable clones can be created from existing VMDKs enabling space efficient provisioning for new VMs. Similarly, snapshots can be created and managed according to policies that support data protection schedules. Finally, offloading the snapshot and cloning processes to the storage controller, which can be triggered from the integrated vCenter console, frees up the virtualization host from this overhead and allows the backup or recovery of VMDK files easily. Integration with virtualization platforms can help storage-at-scale systems to greatly simplify the management of this growing part of IT operations, without compromising control.

Another result of the increased efficiencies of an IT operations-at-scale system is a reduction in the overall cost per GB of capacity. Capacity at scale enables better storage density, which lowers the data center footprint for a given storage system, decreasing power and cooling costs. Performance at scale keeps system 'sprawl' in check and costs down, as it reduces the need to overbuy controller horsepower or add more storage systems to keep up with the performance demands of growing data sets. Storage systems designed for operations-at-scale keeps administrative costs down since they don't add management complexity when capacity grows.

Cost efficiency is also maintained by repurposing existing assets and by eliminating vendor lock-in. Having an integrated storage tiering function helps reduce cost and provide investment protection by allowing third party NAS storage to be incorporated into the storage system.

Data Protection at scale

The total cost of ownership of a storage system must include the entire data protection process. For a storage-at-scale system, this is especially important since its very large capacity can require significant resources for backup, DR, high availability and other related services. To get data offsite, most systems rely on some form of change-based replication. The efficiency of an object-based architecture and active metadata management can reduce the amount of data actually transferred independent of storage capacity. This can speed up replication, drive lower bandwidth requirements, consume less storage capacity and reduce CPU overhead at both the source and target.

This efficient metadata processing also enables the storage-at-scale system to support more snapshots and clones, maintaining data protection effectiveness as the data sets grow. A virtualized architecture allows multiple storage servers to be created and migrated to load balance the system as it grows, and provide the HA failover required by each file system.

In order to support seemingly endless growth in unstructured data many organizations need enterprise storage systems that scale into very high capacities. But more than just expanding into the PB range, these enterprise file storage systems must also maintain efficiency as they scale so that performance doesn't degrade in the process. A storage-at-scale system like BlueArc's, with the flexibility to expand capacity and performance independently, while keeping cost per TB and operational overhead under control may be the answer.

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