STORAGE SWITZERLAND

SAVE AS YOU SCALE CAPACITY

Storage Switzerland, LLC

George Crump, Senior Analyst

An almost universal truth about network attached storage (NAS) is that adding capacity to the system will be a constant need. More data is being generated, it's being kept longer and the size of the individual files are larger than ever. File based data or user data is indeed the dominant area of growth and expected to remain so for the foreseeable future. The growth in these data sets have stretched the systems that store them to their breaking points and have made the cost to scale capacity a top issue with data center managers.

In the legacy environment, file based data growth led to file server sprawl which created the market for NAS appliances. Legacy NAS systems were implemented on similar hardware that ran file servers but with optimized software and operating systems to provide better performance. They could handle more users and more capacity than similarly configured server hardware running a general purpose operating system. This enabled customers to reduce the number of file servers required, often to a single NAS.

The Hard Costs of Scaling Legacy NAS

As data growth accelerated though, extra NAS systems were added to the data center because of limits on system capacity, performance or both. The resulting NAS sprawl decreased storage efficiencies as many small volumes had to be created on individual NAS and the balancing of those volumes managed by the storage team.

Legacy NAS systems or fixed storage systems have a point at which storage can no longer be added to the system. In fact most start to see a significant drop off in performance long before their physically supported spindle counts are reached. The NAS controller can only sustain performance to so many disk spindles. When this limit is reached the path to capacity expansion means the addition of another NAS head. With today's data growth it demands the organization quickly needs more and more NAS heads. Each NAS head represents an obvious capital expense plus the expense of additional management.

There is also the hidden expense of inefficiency. Since volumes cannot expand across NAS heads data and typically whole volumes must be migrated between the various NAS silos to balance out both capacity and performance demands. This inexact process consumes admin time and leads to under subscription since most storage administrators would allocate too much capacity on all the systems instead of too little on one and have to rebalance. It also means that an application may be shut down while the migration occurs, which can also be extremely time consuming.

The Costs of Scaling Legacy Scale-out Storage

The initial solution to NAS sprawl was scale-out NAS which spread the capacity and I/O load across many small servers, called “nodes”, acting collectively as one. Each node had its own storage capacity and I/O bandwidth. When a capacity or performance bottleneck was hit another node was added. While this addressed much of the capacity issues it introduced a new set of scaling issues.

One of those was ‘node sprawl’, the addition of complete NAS storage modules at a steady pace to typically solve a single problem. As stated earlier, when a node is added to the scale-out architecture it brings with it additional capacity, I/O bandwidth, processing power and memory. Typically only one of these four is needed and most often it’s capacity. The result is that nodes sit largely idle because only one resource is needed, but scale-out expansion includes all of the other resources as well. While it’s always nice to have additional headroom, in today's IT data center too much headroom should be viewed as wasted resources and most importantly IT budget dollars.

The impact of node sprawl also has a technical component. While it may be invisible to the user there is a significant level of internode management and communication that occurs under the covers in a scale-out architecture. This can impact performance and add latency to the system, slowing storage I/O response times.

The “add a node” strategy of most scale-out NAS systems also brings other hard costs, aside from the four resources listed above. Sometimes called ‘sheet-metal costs’ these include the chassis, power supplies, cabling and associated physical hardware required to support each node. The challenge is that these are fixed costs included in each node, largely independent of the capacity or level of the other resources that drove the addition of that node in the first place. The other option is to very densely pack the nodes with capacity potentially using 2.5” hard drives. This could lead to 10 or 20TB nodes but not every upgrade needs to be 20TB’s in capacity. Additionally there would then be concern about generating enough performance per node to drive that capacity. In either situation the nodes can easily become out of balance.

Capacity at Scale Storage

The alternative solution to scale-out storage may be an expandable NAS platform that can scale *inside* the system to cost effectively meet capacity demands without incurring other cost and performance challenges. The goal of systems like the BlueArc Titan and Mercury NAS is to provide a single, highly optimized, capacity ‘at scale’, platform that’s able to meet the wide variety of challenges now being thrown at NAS systems. There are two major aspects to capacity at scale. The first is making sure that the system stores data efficiently so that capacity expansion is minimized. The second is making sure that the system scales in a cost effective manner, meaning that capacity is added incrementally, with performance as needed, and may include different tiers of storage from different manufacturers.

Efficiency

Storing data efficiently includes the ability to leverage the storage system’s own intelligence in reducing data growth. A legacy example of this is the ability of NAS systems to create snapshots of data sets and to use those snapshots as secondary copies or for backups. Modern NAS architectures have moved well beyond basic snapshots to include the ability to create baseline images that can be used repeatedly across multiple applications. This technique is often called “cloning” although a second copy is not made of the data, the primary copy is leveraged repeatedly.

The ideal use case of this is in virtualized environments where a single server or desktop image is cloned repeatedly across hundreds and, in the desktop case, thousands of virtual machines. The capacity savings in these cases can be dramatic. The problem is that most systems are limited by how many clones of a volume they can maintain without adversely impacting performance. Modern NAS file systems like BlueArc's are able to manage 16X more clones than the typical legacy NAS. Efficiency is a critical component to curtailing data growth, but at some point NAS systems must add capacity. When they do it's important that they grow in the most cost effective manner possible.

Cost Effective Scaling

Dynamic expansion could be described as scaling only when and where it's needed. This means the effective use of current assets and of new assets as they come online. With the wide range and costs, of storage media available today's modern NAS systems should have the ability to automatically place data on the appropriate tier of storage as needed. For example, it should be able to leverage solid state storage for very active files and then move data down to various tiers of storage as needed, eventually landing on high capacity and cost effective SATA based storage. This should also include the ability to incorporate other NAS devices as destination tiers. Instead of replacing old storage systems, leverage them as a near-line tier of storage.

At scale systems have the ability to add storage processors, network I/O bandwidth and storage I/O bandwidth independently. This means that trays of storage can be added to the system like a legacy NAS but an at scale storage system has the ability to add the storage horsepower needed to support those additional spindles. This also makes it easier to mix storage types in a single system. In today's data center there is often a need for high performance solid state storage, high performance mechanical drive storage and high capacity, cost effective mechanical drive storage. Mixing these different storage types is sometimes difficult in a scale out architecture and again a legacy NAS architecture is left lacking when it comes to maintaining performance to these different drive types. At scale systems let you scale exactly what needs to be scaled (performance, capacity, I/O) when it needs to be scaled, instead of buying all three components at the same time.

The big savings with at-scale storage comes when scaling capacity. Since it can be added to internally as needed, the investment in the physical NAS chassis is made only one time. The additional drive shelves are added fully populated and in much higher densities than the scale out systems, leading to better utilization of drive capacity. This architecture also saves potentially the most sought after resource in the data center, floor space. At scale systems are able to very densely package themselves into the existing data center.

Saving as you scale capacity brings a number of benefits. First, there is a natural hard dollar savings provided by the at-scale architecture, compared with the 'sprawl' of legacy NAS systems and scale-out storage. Second, storage efficiencies are increased through the use of storage tiering, snapshots and cloning, plus the fact that volumes are not spread out across multiple NAS heads. And finally, data center floor space is saved thanks to the ability of at-scale systems to densely package capacity, storage processing and storage bandwidth.

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