STORAGE SWITZERLAND

ARCHITECTING A SCALE-OUT DESIGN FOR SCALE-UP STORAGE

Storage Switzerland, LLC

George Crump, Senior Analyst

Ideally, you never want to outgrow your storage system. You'd like it to grow with the demands of the data center both in terms of performance and capacity. If the storage system can't scale, then being forced to buy a replacement or added system is an expensive next step.

As a result the scalability of a storage system is a critical point of evaluation during the selection process. Vendors will use a variety of methods to meet that scaling requirement, including highly scalable single systems (scale-up storage), array based scale-out storage systems and scale-up storage systems that scale out with the addition of software or hardware.

Scaling Should Be Part of the Initial Purchase

The basic building block of any storage infrastructure is the storage controller or storage head. This is the brains of the operation. That controller handles the front end I/O from the server and the back end I/O to the storage. Ideally, the capabilities of this initial building block should be enough to support the storage performance and capacity demands of the organization for at least a few years. While this article discusses several ways to overcome when and if the limits of that initial block are reached, it should not be something that has to be addressed within the first year or so of purchase.

Sophisticated array-based or software-based scale-out storage configurations may not be needed for many data centers. Modern day storage systems have very wide initial performance and capacity scaling ranges. They can start at configurations ideal for small enterprises and expand to support the demands of the very large data center. For example, the Nexsan E60 storage system has a powerful enough controller to scale from minimal capacity of 60TB to 180TB while supporting 25,000 IOPS at full capacity (a realistic blend of 75% reads with 25% writes of 4k blocks). It can then be doubled in capacity by adding an expansion shelf for a total of 360TB.

However, at some point there may be an occasion where performance or capacity of these systems has been reached. This is where the array-based scale-out storage systems that Storage Switzerland discussed in its previous article, The Advantage of Scale Up Storage, may be considered as an alternative by storage managers. The challenge with array-based scale-out storage is being able to start small while still delivering acceptable performance and then actually scale large enough when the demands of the data center warrant it. While in theory some of these systems claim to scale to 100 or more nodes, the data center has to consider the space that large numbers of nodes would consume and the power they would require, not to mention the complexity of cabling 100 nodes.

By comparison, traditional scale-up storage handles the first situation very well. These systems can start small while delivering excellent performance and can scale very effectively to a point that meets many data centers' potential requirements. For those that do need to expand outside their initial system there are several options available. Many of these can allow the data center to leverage the benefits of a scale-up architecture while still getting the same growth capabilities of an array-based scale-out architecture.

Scale-in-Place

Designing a scale-out storage architecture typically implies using additional software and hardware to integrate multiple storage controllers so that they can be managed by a single interface. There are many ways to accomplish this without having to buy into an array-based scale-out storage system. A first option to consider is to leverage multiple stand-alone units and indeed manage them separately. Many storage systems today include an interface that makes it very easy to switch between managing two separate controllers and many users will find that managing two or three is not any more difficult than managing one.

There is also a unique advantage to this method of scaling storage. If the storage system has a power management function like Nexsan AutoMAID® energy-saving technology, a new system can be brought in and have only new and active volumes placed on it. This gives the full storage capacity and performance capability of the new system to the most active or newer volumes. Then, since the remaining volumes on the older system are typically accessed less frequently, MAID levels can be set so that the system can spin down to a more power efficient state. Scale-in-place allows the data center to benefit by improving performance with the new system and save on power and cooling on the legacy system.

Scale with Software (hypervisor, foundation, cluster software, cloud software)

At some point the management of multiple storage systems does become a problem and a scale-out architecture needs to be considered. This does not necessarily mean replacing the scale-up storage system with an array-based scale-out storage system.

Doing so would not only mean losing the initial investment in the scale-up storage system as well as a potentially painful migration process, it would also mean losing the advantages of a scale-out system.

Appliance-based scale-out storage

One method to scaling out an existing scale-up storage system is to use an appliance or software running on a server that will virtualize multiple storage systems and enable them to be managed as one, leveraging the power of all the controllers.

For most environments the value in this type of storage virtualization is providing a single point of management. But for environments that need incredibly large storage capacity or more likely, massive amounts of storage controller processing, some systems have the ability to create a volume from multiple storage systems leveraging the storage controller of each into a raw aggregate that delivers an unprecedented level of performance.

The advantages of this solution are that you get the benefits of scale-up storage in the form of highly dense, very powerful storage controllers, but now combined into a single entity. You also get the advantage of flexibility. For example, you could mix and match systems as you see fit, something that can help keep the vendors honest. Storage vendors that make great hardware that's also reliable and efficient typically have no issue being placed into an environment where they need to earn your business with each successive storage purchase.

The Storage Hypervisor

Another method to scale-out scale-up storage is to install the scaling software within the hypervisor or at the hypervisor layer. Obviously ideal for server virtualized environments, these solutions bring a natural scaling capability to the environment as we describe in our video Storage Hypervisor. As more virtual hosts are brought online each one shares in the scale-out storage responsibility.

As with the appliance approach mentioned above the combined aggregate power of the storage controllers on each individual storage system can now be leveraged as one. Once again this brings flexibility and freedom of choice to the storage manager.

Storage Cluster or Cloud Software

Another option is to use storage clustering or cloud storage software to bring scale-out capabilities to scale-up storage. Using software from providers like Veritas or Red Hat (thanks to their acquisition of Gluster) a data center can build its own cloud-like storage infrastructure.

Typically, these systems leverage individual servers with individual storage systems attached to them. The cloud or cluster storage software is installed on the servers making each server a node in a global storage cloud. The storage manager, and of course, the users interact with this cloud as if it was one entity. The software behind the scenes manages which node is most appropriate to store which data and the redundant copy of data.

The systems are typically available in either a loosely coupled cluster or a tightly coupled cluster. In a loosely coupled configuration data is held discreetly on an individual node and then typically copied to another node or nodes for redundancy. In a tightly coupled cluster, data is typically segmented or 'chunked' and then wide striped across all or some number of nodes in the cluster.

Both methods can provide tremendous performance benefits since the clustering software typically balances I/O performance among the member nodes. Loosely coupled clusters generally provide more flexibility in the types of servers that can act as nodes as well as the storage attached to those nodes. But performance per data file is limited to the performance of an individual node. Tightly coupled clusters tend to be slightly more restrictive in the types of servers that can act as nodes and the mix of storage systems they support. But because data is segmented, all the nodes are responsible for delivering data when requested. Therefore, the performance is the aggregate of all nodes - not the maximum capabilities of an individual node.

The Advantages of Scaling Later

Scaling later has some distinct advantages over scale-out storage. First, you get to start with scale-up storage, which is typically a smaller upfront investment and often easier to implement, especially from a networking perspective. Later, one of the above scale-out methods can be added to the scale-up system and provide much the same benefits of an array-based, scale-out storage system.

More importantly you gain the benefits of no vendor lock-in because storage systems can be selected based on the task at hand instead of being bound by an initial decision that may have been made years ago. Secondly, if a storage system like that of Nexsan is chosen, you also eliminate redundant expense since you don't have to re-buy features like snapshots, replication, and cloning. Finally, you should also see more efficient scaling where instead of the goal being to add 100 nodes that consume data center floor space, you had 10 nodes that deliver more power and use significantly less floor space.