

Virtualizing NAS for Green Storage

Virtualized NAS clusters can save you a cluster of power and money

April 21, 2008

By George Crump, Founder and President, [Storage Switzerland](#), April 21, 2008, 12:35 PM

Power generation and distribution are significant cost considerations when building a new data center. If you can use less space and do so with less power, you can reduce a key budget hog.

This isn't just a matter of immediate concern. Many organizations are running out of the ability to provide power long before they are running out of the ability to pay for that power. As a result, it's vital to be "green" when planning a new or improved data center for the future.

That said, many green initiatives are stuck in the forward position. While most green IT has moved from something to talk about to the point where project plans are being created, in most cases, these plans are elaborate and will take months, if not years, to fully implement. Even the commonly accepted process of moving to a virtualized server environment is a multi-month project.

These timelines stand in contrast to the fact that many data centers are running out of power *now*. What can be done that does not require months of planning and implementation?

We think there are a few simple steps involving network-attached storage that can provide immediate power savings while also providing additional benefits to the enterprise.

File services consolidation

At the top of the list of immediately accessible green storage options is file server consolidation by implementing a NAS. Simply reducing the number of file servers required by the enterprise can cut power costs.

However, more than just NAS is required. Most data centers already have purchased a NAS product to consolidate file services. The weakness of traditional NAS solutions is that they operate with individual standalone storage controllers that cannot be virtualized. Thus, as the need for more I/O capability or storage capacity grows, the traditional setup calls for an organization to add additional separate controllers. This leads to managing independent islands of NAS systems. So over the course of time, that initial effort of file server consolidation leads to distributed NAS systems, in essence recreating the same problem that

NAS was meant to solve.

The multiplication of NAS systems is a real-world concern, especially considering that unstructured data (file data) is the fastest growing data set in the enterprise today. Consolidating to a traditional NAS is then not a long-term strategy but merely a Band-Aid, because of its lack of ability to scale in a way that can be practically managed. Traditional NAS should not be considered as a long-term solution to lowering power consumption.

Virtualized NAS cluster

If a traditional NAS system is not the solution to green file services, something different is needed -- a virtualized NAS cluster. This is a grid of NAS controllers that provides virtual file servers within individual NAS nodes. This technique allows for scaling across NAS nodes to address performance, scalability, and redundancy concerns while at the same time scaling down to provide granularity. Virtualized NAS clusters address the growth vectors that most data centers will experience.

In a virtualized NAS cluster, now offered by vendors like Isilon and ONStor, with NetApp planning an OnTap GX product this year, each virtualized NAS controller is a node in a grid that presents a single point of management and control. When additional I/O performance is needed because of user growth or increased data loads, another virtualized node is added to the cluster. That new node is now managed with the other nodes in the grid; no separate management point is required. If more capacity is required, more physical capacity can be added to the cluster.

The independent scaling of NAS nodes and physical storage provides maximum flexibility, allowing specific growth where the need is the greatest. This results in power savings. In contrast, traditional NAS calls for the purchase of capacity and performance together when only one is needed, thus wasting power.

Things to look for

The virtualized NAS cluster should be able to mix different node types. This allows for file systems that require greater performance to be assigned to nodes in the cluster that can deliver more performance. Slower nodes can be leveraged in the case of a failure.

Being able to cluster different types of nodes also provides investment protection. With virtualized NAS clusters, older technology can be retained while moving some nodes to faster, more reliable, and more power-efficient platforms as those become available.

Support of mixed virtualized nodes also helps eliminate full-scale migration, a problem for traditional NAS. With mixed nodes, manageable migrations can happen over time – or immediately – as the more advanced nodes become available. As the NAS cluster grows, this is a key requirement. When dealing with data stores measured in petabytes, it is almost impossible to move all the data to a new platform at once.

A common challenge with simply building a large NAS or file server is that there is a lack of granularity and as a result, most of the NAS is idle. With a virtualized NAS cluster, not only

can adding a node increase the performance, the services the nodes provide can be virtualized. Each node in the grid can run multiple virtual file servers if needed. This not only allows for greater granularity and for security to support department-level requirements within the organization, it also significantly helps with the initial migration to the NAS.

This brings up another benefit: Since a virtualized file services architecture can sub-divide into multiple virtual file servers, each file server to be replaced by NAS can be virtualized into that environment. The users see no change in the way they interact with the system. As a result, migration to the new, more power-efficient virtualized environment is quick and seamless. In contrast, if consolidation of file servers requires redirection of users to a single new mount point, it will create user confusion, take time, and be very disruptive to the environment.

Virtualized NAS clusters enable the move of file systems to new, faster nodes or it enables the movement of an old file system to less expensive and higher capacity SATA drives. Power requirements are further driven down by storing more data in less space, yet that data is still easily accessible in case it is needed again.

A component of any good virtualized NAS cluster is a global file system. The global file system brings all the components together by allowing specific file servers, file systems, or folders within the file systems to be moved to different nodes or types of disk without user disruption.

Cost savings

As an example of how these power requirements deliver cost savings in a data center, look at a typical data center with 1,000 users accessing 24 servers storing 6 Tbytes of actual used storage on 20 Tbytes of total capacity.

The virtualized NAS Cluster can deliver the same performance with 4 nodes and 10 Tbytes of capacity supporting the same number of users. The space used by the 24 servers is over 48U, the space utilization of the virtualized NAS cluster is 8U, and power utilization drops from 25,600 watts to 1,680 watts, a savings of 93 percent. Other savings result from eliminating free excess capacity, as well as enabling a more flexible time-responsive environment. For example, storage can be added automatically and resources can be reallocated in seconds.

Bottom line? A virtualized NAS cluster not only solves power efficiency issues, it also creates a more flexible file services environment for users, making it a "go to" technology for the data center.