White Paper

Five Things to Consider Before Upgrading Your NAS

Executive Summary

Businesses looking to boost the performance of their NAS environments typically look at a variety of potential upgrades to solve their problem, including adding more hard disk drives (HDDs), short-stroking HDDs, installing Flash PCI cards or solid state drives (SSDs), or upgrading to higher performance and much more expensive NAS controllers. However, these solutions can increase costs unnecessarily, offer minimal performance increases, and are short-term solutions that cannot scale to keep pace with a company's need to serve new clients and applications. This paper examines the top five issues to consider before spending scarce capital resources on new NAS upgrades.

The True Cost of HDDs

In today's homogeneous-media NAS architectures, storage administrators are invariably asked to add HDDs to increase overall capacity or system performance. Beyond the acquisition cost of the hard disk media, administrators need to carefully consider the total cost of adding HDDs to an existing or newly planned NAS infrastructure. Beyond the acquisition cost of enterprise-class Fibre Channel (FC) drives, you also have to factor in the hidden costs of power and cooling, the necessary rack space, and data-center floor space. Furthermore, because HDDs are inefficient at providing IOPS, adding expensive HDDs to achieve performance goals is not a good investment from a cost-performance metric. The value proposition becomes even worse when companies overprovision capacity and short-stroke drives in attempts to achieve even higher performance.

Instead of trying to extend HDD technology, NAS users will get a better return on their storage investment by moving from a homogeneous solution to a tiered NAS solution. Adding intelligent storage tiering that makes the best use of newer SSDs produces a dramatic gain in IOPS performance, with much lower consumption of power, cooling, rack space, and data-center floor space. This approach eliminates the expense of adding potentially hundreds of enterprise HDDs to the NAS environment, recovers the capacity on existing HDDs that was previously unavailable due to short-stroking, and solves NAS performance issues while enabling IT to leverage low-speed (and low-cost) SATA drives for capacity requirements.

Inefficient Use of Flash Storage

Integrating Flash storage into NAS systems is a new method of improving performance that is being promoted by vendors of traditional NAS. List pricing for Flash in a NAS system from one of the leading vendors runs from $170/GB to $300/GB. This compares with a list price of roughly $2/GB for SATA storage. At these prices, customers need to be using Flash very efficiently. Unfortunately, vendors are not always making this possible.

Typically, Flash is added to a NAS system in one of three ways: as a PCI card, an SSD array, or a caching appliance. Let's look at the leading vendor in each category and the efficiency of their approach.

AVERE

Part number 0254-002-0171

NetApp offers Flash as a PCI card they call PAM. Up to five PAM cards with a total of 2.5TB of Flash can be installed in some NetApp controllers and pricing is in the $170-200/GB range. PAM is inefficient for two reasons. First, PAM is read-only, so customers still need many hard drives to handle the write workload. Second, PAM helps the read performance only on the controller in which it is installed. This means you might add up to five PAM cards to every NetApp controller in your environment. This approach is expensive and highly inefficient in environments with multiple NetApp systems because Flash is added to individual storage "silos" to handle the peak load on the silo but is underutilized much of the time.

EMC offers Flash as an SSD array. Up 16 SSDs and 6.4TB of Flash are supported per array and pricing is in the $200-300/GB range, making a fully populated array extremely expensive. Beyond the price, EMC's SSD arrays are inefficient because data movement to Flash is slow and not granular. Data movement is triggered by a policy engine that measures data activity across long periods (e.g. days) and cannot immediately respond to a hot application. Additionally, data is moved in large volume-level chunks, meaning lots of cold data is using expensive Flash storage alongside the hot data.

Avere offers Flash in an appliance that sits in front of NAS systems from other vendors. The Avere FXT 2700 has 512GB of Flash per appliance, scales to 25 appliances (13TB) per cluster, and accelerates both read and write workloads. The Avere architecture is highly efficient since it provides a consolidated Flash layer that is shared by all of the backend NAS systems. This allows customers to add the right amount of Flash to their NAS infrastructure to deliver the required aggregate performance. In addition, the FXT 2700 makes the most efficient use of the Flash storage by moving data in real time and at the finest level of granularity possible, blocks within files.

Disruptive Upgrades

With traditional NAS, controller upgrades are part of the typical lifecycle of the system. Controllers are purchased to meet current performance requirements, but future performance requirements are not factored in because that would mean spending more money. A year or two down the road, the controllers cannot keep up with increased performance requirements and something must be done to get more performance for the data.

At this point controller upgrades are the typical course of action. Controller upgrades are so common that people have become numb to the pain. Let's describe the problems that come with this and propose a better way.

Controller upgrades with a traditional NAS system involve the following steps:

1. Purchase higher-performance controllers (typically a model or two up the NAS vendor's product line, if a higher model exists)
2. Purchase new disk shelves or Flash-based PCI cards, required to get more performance out of the new controllers
3. Purchase new licenses for all software (e.g. NFS, CIFS, mirroring) at the higher price tier of the new controllers
4. Take the original NAS system offline
5. Remove the old controllers
6. Install the new controllers
7. Add the new disk shelves
8. Bring the upgraded system back online

The above process is expensive, disruptive, and requires more disks, power, and space even if no additional storage capacity is required. As an alternative, consider boosting the performance of your existing NAS by breaking your NAS architecture into two stages by adding a performance-oriented automated tiering layer in front of your existing controllers.

Future Scalability

If you're at all concerned with the scalability of your infrastructure when considering upgrades, you should know that adding new controllers, more high-speed drives, or Flash modules to a NAS installation to improve performance is a short-term solution at best. It's only a matter of time before application demands once again outstrip the NAS infrastructure's ability to scale performance and you're back to ripping out old gear and replacing it with new.

In contrast, system scalability is built into Avere's two-stage NAS architecture. As more clients and new applications are added to the mix (requiring higher IOPS performance), an Avere FXT cluster can be easily expanded with the non-disruptive addition of new nodes. Up to 25 FXT nodes can be added to a cluster, delivering plenty of horsepower without having to touch any other devices already in place. And because the Avere FXT cluster can serve multiple storage servers, there is no need to

add Flash to each and every controller – the Avere cluster becomes an extensible fast media layer in front of all of them, providing high performance to hot spots without overprovisioning.

Ease of Management

Manageability is another hidden cost to upgrading an existing NAS infrastructure. With falling prices and improved durability making new storage media such as Flash SSDs widely available to boost application performance, many companies are tempted to install Flash at tier zero and expect that it will solve their performance problems, albeit it at a relatively high cost.

However, installing fast-access storage media solves only part of the problem. IT then has to figure out which applications are best served by the new tier of storage, often having to become an expert in the latest storage media read and write rates and application QoS schemes in order to optimize the utilization of the more costly storage media. In comparison, an Avere FXT cluster has the intelligence to dynamically allocate data to the appropriate storage tier and media based upon both data and access characteristics, which balances the cost/performance equation with no administrative overhead.

Summary

Before upgrading your NAS infrastructure, do an analysis of the true cost of adding or short-stroking performance-oriented hard drives, filling controllers with Flash cards, or upgrading to the next level of NAS controllers. You may find out that these solutions are expensive, inefficient, or merely a short-term solution that limits future growth. Consider implementing a tiered NAS approach like that offered by the Avere FXT series of appliances. Avere's solutions boost the performance of all NAS applications by accelerating read, write, and metadata operations without the need to add more disks to your NAS system. FXT appliances scale from 1 to 25 nodes per cluster to match your initial needs and gracefully scale as your needs grow. For capacity, simply add SATA drives to your existing NAS system. FXT appliances are simple to install in existing environments and require no changes to existing applications, clients, NAS systems, or data retention procedure such as backup and mirroring. By carefully considering the best path for NAS upgrades, you can improve application performance while saving your company significant cost.

Avere's solutions boost the performance of all NAS applications by accelerating read, write, and metadata operations without the need to add more disks to your NAS system. FXT appliances scale from 1 to 25 nodes per cluster to match your initial needs and gracefully scale as your needs grow.

AVERE

Avere Systems, Inc. | 5000 McKnight Road, Pittsburgh, PA 15237 | 412-635-7170 | Fax: 412-635-7171 | averesystems.com