



HP Scalable NAS

Product Insight

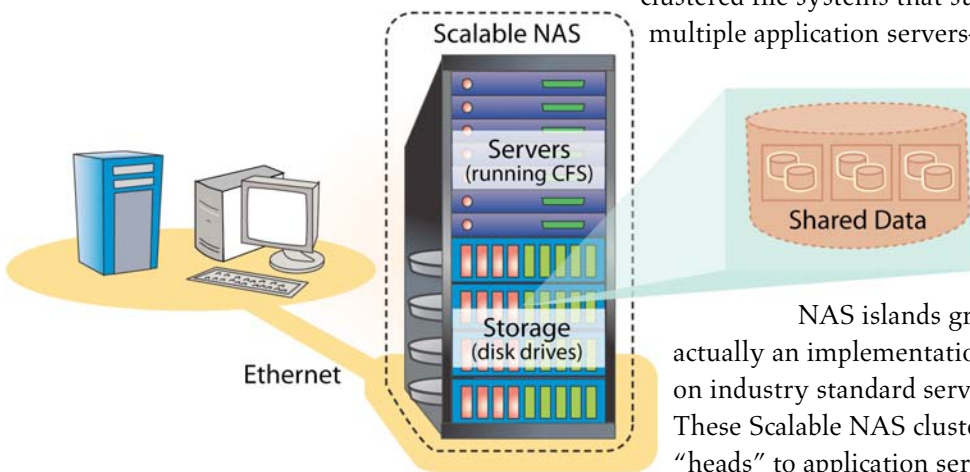
John Webster
15 August 2008

Licensed to Hewlett Packard Company for web posting. Do not reproduce. All opinions and conclusions herein represent the independent perspective of Illuminata and its analysts.

IT administrators, new to networked storage concepts and implementations, normally ask: What's better, networked attached storage (NAS) or storage area networks (SAN)? The answer has always been: It depends—on the application mix, the existing environment, performance and availability requirements, anticipated capacity growth, and most important of all, available budget. Early adopters of NAS discovered that they could use the Ethernet infrastructure they were familiar with to build networked storage pools for significantly less money than a SAN with equivalent capacity, and could still support transaction-oriented applications.

Soon after the initial euphoria, NAS was dogged by an inability to scale effectively as IT saw relentless demand for more and more capacity. The saying: "You'll love your first, but you'll hate your tenth," became the common way to describe the experience of storage administrators who, over time, were continually adding NAS platforms simply to keep up with capacity demand. Ironically, these users who turned to NAS to consolidate and integrate storage islands trapped on dedicated servers were recreating the same problem they were trying to solve, only on a grander scale, by adding NAS platforms that couldn't be consolidated and integrated.

Vendors have long been aware of the NAS scalability problem, but it's taken them years to bring solutions to market. Not a minute too soon for storage administrators, scalable NAS—based on the global namespace concept and clustered file systems that support true data sharing across multiple application servers—is now showing up in production IT environments.



HP's Scalable NAS is one approach that can successfully address the administrative problems that escalate as the number of

NAS islands grows. HP Scalable NAS is actually an implementation of HP's clustered file system on industry standard servers attached to SAN storage. These Scalable NAS clusters present themselves as NAS "heads" to application servers or clients. Since the number of servers (NAS "heads") in a Scalable NAS cluster scales to 16, and the size of the sharable storage pool scales to 2 PB, Scalable NAS is HP's version of the next step in the evolutionary path of NAS.

In this paper, we tap into the experience of storage administrators who currently manage HP Scalable NAS installations. What characteristics drove them to embark on these projects? What are their experiences? And how does it compare to more traditional NAS approaches?

What Motivated IT Administrators

We recently spoke to a number of HP NAS users from a range of production and application settings, including digital video, Web 2.0 content delivery, healthcare, and messaging. These users tended to fall into two categories:

The “We Need Something Big” Group – Those who were building new networked storage platforms to support new applications (i.e. “greenfield” storage projects), and realized that scalability was a critical requirement.

The “We’re About to Hit a Wall” Group – Those who could see a certain futility in continuing down the storage path they had been traveling and felt the need to depart significantly from that path.

We Need Something Big (Maybe Really Big)

As users contemplate the requirements for a new application or new business opportunity, they realize that the approaches they are used to taking won’t work. They won’t scale, won’t handle the projected load, or can’t be managed by existing staff. The mountain to climb is too high for the solutions they’ve already experienced. They realize that they need to take a new approach, possibly a radically new approach, capable of scaling to new heights without adding management overhead.

Customers in this category included those in media and entertainment and Web 2.0 content serving. For example, one customer was transitioning from standard video processing to high definition (HD), a move that would require 10x the amount of processing and storage capacity. Simply building a bigger version of their current architecture wasn’t going to work in this case. They needed to take a radically more scalable approach.

We’re About to Hit a Wall (If We Haven’t Already)

As an application platform matures, it can run out of capacity. It is often the case that IT architects misjudge platform scaling requirements. Capacity is generally used up quicker than anticipated. Here, capacity is defined in terms of:

1. Physical capacity (e.g. compute power, storage disk space)
2. Network bandwidth needed for an acceptable QoS level as experienced by application users
3. IT staff resources required to adequately support the platform

Given the reported average storage growth rates in the range of 80% per year that continue apace, the storage that supports an application platform can easily reach the top end of its ability to scale. Something must be done before the application breaks.

The resulting attitude among storage administrators is often: Let’s do it right this time. People realize that band-aid solutions won’t get them out of the hole that they’re about to fall into. They also realize that they need to take a radically new approach that places them back on the lower end of the manageable capacity scale and gives them lots of room to grow.

What HP’s Scalable, Multi-node NAS?

The users we interviewed from both groups chose HP Scalable NAS. While the approach was different from what they were familiar with—and required a new learning curve—they were willing to devote the necessary time and attention to get something in place that they believed would continuously support their required applications despite an anticipated demand for new storage capacity, even when that was in excess of 100% or more per year.

HP’s Scalable NAS presents itself to the application server or clients as any traditional NAS subsystem would. However, the similarities basically end there. Where one normally sees a NAS “head” or controller, Scalable NAS offers a cluster of up to

sixteen server nodes. The cluster offers full fault-tolerant redundancy, and is attached via an underlying SAN to disk storage arrays. A single HP Scalable NAS cluster can support:


- Up to 16 nodes
- Up to 128 TB of data per file system for Linux (Up to 32 TB filesystem size on x64 architectures only with 8K block size for Windows)
- Up to 256 file systems per cluster
- Up to 2 PB of storage per cluster

The underlying file structure (managed by HP Cluster File System) is fully sharable by all application servers and clients attached to the cluster, meaning that individual files can be concurrently accessed and written to by more than one application server while file consistency is maintained. Processing and storage capacity is added non-disruptively, and performance scales linearly as capacity is added. Finally, and perhaps most important for those familiar with the “You’ll love your first, you’ll hate your tenth” scenario, operational management of the entire scalable NAS

cluster is consolidated and focused onto one platform.

HP offers a number of Scalable NAS implementation scenarios that range from providing the essential elements (some assembly required) to a soup-to-nuts, pre-integrated and installed solution (no assembly required). For example, a customer could introduce Scalable NAS to an existing network, using re-dedicated servers for NAS controllers and existing storage arrays. However, this approach would likely require a significant commitment in staff time to learn, test, and install, and some amount of application downtime. On the other hand, HP could bring in a pre-integrated and tested configuration that included servers, disk arrays, and software; install and test the configuration, then hand the “keys” over to the customer.

To that end, HP Scalable NAS is offered in four basic configurations that differ in the number of components, ease of installation, and price. (See table.) Each of the selections will yield a single, logical NAS cluster with single global namespace that is fully fault tolerant.



HP-Scalable NAS File Serving Software (software only)

- Supports Windows 2003 Server, SuSE Linux, or Red Hat Linux
- Installation plus NAS gateway and storage configuration service using customer-supplied hardware is an extra option. Currently HP supports integration with HP EVA, MSA, and exDS9100 storage arrays.

Enterprise File Services Clustered Gateway (EFS-CGW)

- Includes HP Scalable NAS File Serving Software and redundant HP DL380 servers
- Windows Storage Server (64bit) or SuSE Linux (64bit) pre-installed on the DL380 servers
- Pre-configuration, subsystem installation and integration with customer’s existing storage is included.

Enterprise Virtual Array File Services (EVA-FS)

- EFS-CGW (as above)
- HP EVA storage
- Includes FC switches and IP switch
- Factory-integrated bundle
- Installation and configuration service included

Extreme Data Storage (ExDS9100)

- 4-16 ESF-CGW with 246 TB to 820 TB of storage
- Factory-integrated bundle with single management interface

Several additional options are available. These include:

- A “stretch” cluster capability that adds HP-CA for synchronous storage replication is available for installations requiring inclusion within the organization’s disaster recovery plan.
- HP Data Protector for backup of the entire cluster
- TrendMicro antivirus software

What We Heard

Each of the customers we interviewed reported variations in their deployments. Nevertheless, some common experiences emerged from our discussions:

Performance is deterministic as capacity is added. The addition of compute power and/or storage bandwidth yields a predictable increase in performance, allowing for orderly and more manageable growth. This perception was particularly true for those customers who installed Scalable NAS when they felt they were about to “hit the wall” with their existing storage architecture.

Single console management allowed administrators to apply upgrades and patches across the cluster once as opposed to doing so to each server in the cluster. This was appreciated by all. Performing repetitive operational tasks on multiple servers and storage devices keeps administrators unnecessarily occupied while other more meaningful jobs and/or applications users are kept waiting in the queue. Repetitive iterations of the same task without single console management is also error-prone.

Time devoted to problem resolution was significantly reduced. One customer from the “hit the wall” group reported that it could take a team of five administrators hours to resolve a problem in their previous environment. In fact, it was downtime that induced them to look for a radically different architecture.

Focused management was commonly cited by the group who chose HP Scalable NAS for a new

application. IT staff resources are generally burdened if not over-burdened to begin with. Therefore, the ability to add cluster nodes and disk capacity to a single storage pool without disruption was significant.

Eliminating the need for file replication across servers was commonly appreciated by users in media and entertainment. In this sector it is common to find application users passing files back and forth from server to server as they process a video file, for example, through various processes to completion. Here, true data sharing eliminates this constant file thrashing that can greatly elongate project delivery cycles.

HP Scalable NAS and ROI

Generally speaking, the most significant ROI benefits resulting from implementing HP Scalable NAS came from gains in application user productivity. For example, a video processing customer who installed Scalable NAS to streamline the production process saw an immediate increase in monthly revenue simply because they were able to get more jobs per month out the door and invoiced. The administrative staff that was spending hours to troubleshoot an Oracle application outage saw an immediate increase in application user productivity because availability and performance increased significantly. Such benefits can be quantified and substantial. We encourage potential Scalable NAS customers to explore this type of ROI analysis; the results could facilitate the internal acquisition process.

We also encountered another type of ROI when we interviewed a Scalable NAS user in healthcare. Here, a digital medical imaging system was running out of capacity. As a result, performance was steadily degrading. The entire imaging archive was migrated to Scalable NAS. Performance improvements allowed patients to be diagnosed and treated more quickly, improving patient care and quality of life.

Conclusion

As an IT administrator, if you are experiencing of the following issues, you may be in line to take a different architectural approach to your storage environment:

- Application performance is degrading as you add capacity.
- Adding capacity, installing new software revisions, and applying patches requires downtime.
- Operations management is distributed across multiple platforms in a way that forces administrators to perform the same tasks repeatedly.
- Data that could and should be shared, can't be shared among servers, application processes, and application users.

While HP Scalable NAS is not the only available clustered file system-based NAS product currently available, it is one that is fully functional across a wide range of application scenarios, and is fully integrated, supported, and available from one

vendor. It was chosen by the IT administrators we interviewed over others specifically because, for them, it does a better job of addressing these issues. However, we caution potential customers to fully evaluate the progression of Scalable NAS solution offerings from a “learning curve” perspective. NAS based on a clustered file system from any vendor will require an education. HP has also learned from its customers and we commend HP for offering a range of options that include system integration, installation, and on-site testing.

We believe that, as data center architects take advantage of high-bandwidth Ethernet infrastructure in increasing numbers, and continue to virtualize their server environments, NAS—being both Ethernet-based and virtualized from the start—will receive increasing scrutiny and acceptance. Now that the scalability, availability, performance, and manageability issues are being addressed with clustered file system architectures like HP's Scalable NAS, they can feel comfortable with the new NAS as a platform for critical application storage.



Through subscription research, advisory services, speaking engagements, strategic planning, product selection assistance, and custom research, Illuminata helps enterprises and service providers establish successful information technology.