NAS VIRTUALIZATION

Selecting The Product That Fits Your Needs

White Paper

January 2006

by Farid Neema
Peripheral Concepts, Inc.

NAS Virtualization

NAS virtualization is a software-based solution that, in its broadest definition, lets you manage a heterogeneous environment including stand alone and clustered servers, workstations and PCs as one storage pool that can be managed, monitored and administered from any console. When one device reaches its limit, the NAS storage virtualization software automatically starts allocating storage from that device to another NAS device or file server in the same pool . Better yet, the more sophisticated systems stripe files across RAID systems connected to different servers and re-stripe them transparently when new RAID systems are added.

1- WHO NEEDS IT?

A NAS (Network Attached Storage) system is a single-purpose device designed to be easy to purchase, implement and use. It combines a file server with the storage array in a single, integrated appliance, also referred to as a filer. A NAS system presents data in the form of files, allowing consolidation of storage that can be shared by other servers on the network. The simplicity of deployment of NAS as a solution to file services has been very attractive to IT managers regardless of the size of the site. But while managing a few NAS devices is easy, managing many has proven to be challenging and costly.

A single file system that expands to the limits of a server's physical capacity requires a time-consuming and complex upgrade to a larger server with more capacity. An easier solution is simply to deploy another server or NAS device, which itself eventually becomes a "data island" when the next NAS device is added to the line. When the addition of capacity is multiplied by a large number of NAS devices, administrators must manage each server or NAS device, each with its own file system and each requiring independent manual attention for both file-system and capacity management issues.

This growing server count not only proves costly and inefficient, but it also puts a significant strain on the IT manager's ability to adequately backup and protect critical data. Furthermore, most filers support only a single type of storage, making it difficult to tier different service levels for different business requirements. All this significantly impacts what is referred to as The Total Cost of Ownership (TCO)

Finding ways to hide complexity, automate tedious tasks, streamline administration, and still satisfy the requirements of high performance and low TCO is the promise of NAS virtualization.

2- IMPLEMENTING VIRTUALIZATION

Traditionally, a Storage Virtualization Layer refers to a level of abstraction implemented in software that provides address mapping between a number of real entities and another number of virtual entities. Applied to disk drives, virtualization divides available storage into virtual volumes. Virtual volumes are used by an Operating System as if they were such disk drives. The Storage Virtualization Layer redirects I/O requests made against a virtual disk to blocks in real storage. The system can move physical blocks and update the virtual-to-real mappings at any time. Rather than having to administer data by working with each disk in a disk farm, virtualization allows IT managers to deal with multi-disk storage as a single logical entity.

You can implement or incorporate layers of virtualization functions at several levels and with several degrees of sophistication. One of the basic levels is found in volume managers, software utilities that create logical volume groups out of physical disk drives. RAID is a special case of disk volume management that provides virtualization. Most operating systems have some level of volume management built in, or layered, using third-party products. Most databases use an

underlying file system to provide some level of virtualization so applications can address a space rather than a physical disk. Network attached storage (NAS) systems also provide a more advanced level of volume management and virtualization level.

Storage virtualization implementations use different architectures and can be implemented in many ways along the I/O path. Most current implementations are in the **host**, in the **storage** system, or somewhere in the **network fabric**.

The virtualization function is implemented **in-band** or **out-of-band**. An *in-band* implementation intercepts the data and controls access to the virtual pool of storage that it manages. RAiD and NAS fall in the in-band virtualization category. The appliance can become a bottleneck and can constitute a single point of failure. Multiple-server virtualization alleviates the performance concerns of in-band virtualization and allows for server fail-over to occur when one of the appliances is down.

In an *Out-of-band* implementation Control functions and virtual-to-physical mappings are handled by a separate entity, while the data flows directly between the hosts and the storage devices. The drawback to out-of-band approach is that it does not provide read/write access to active data. This means that the out-of-band cannot handle open files to ensure continuous access

NAS Virtualization

The concept of virtualization is applied in NAS to aggregate multiple NAS appliances, so they can function as a single managed device. The aggregation is enabled by a new file system that can either replace the native file systems or layer on the top of existing file systems.

A **Global File system** allows files to keep their namespaces. A client needs only configure access to one file system to get access to any of the data in the system. All files that have been moved will still appear to be in their original directory and path.

The most common implementations for NAS virtualization are

- **NAS aggregation appliances**

They front-end traditional NAS servers. They can either aggregate homogeneous servers from a single vendor, or heterogeneous servers from a number of different vendors. This approach virtualizes all storage resources across multiple NAS servers. By residing between the clients and existing NAS storage, these servers virtualize the storage to the clients. This implementation adds one layer of management with its processing latency, negatively affecting performance.

- **NAS gateway**

A NAS gateway front-ends a SAN fabric allowing storage to be managed in a pool. Some implementations manage files with storage capacity assigned to file systems across multiple devices

- **Distributed Server Virtualization**

This approach presents no additional layer above each server. It eliminates the need for a separate virtualization appliance. It also alleviates the performance concern of in-band virtualization.. By adding pooled NAS servers rather than adding space behind a single server, enterprises can scale performance and capacity at the same time. This approach effectively makes each server's capacity part of a seamless pool.

This virtualization is achieved through a **Distribute File System** (DFS). Using Ethernet as a networking protocol between nodes, a DFS allows a single file system to span across all nodes in

the DFS cluster, effectively creating a unified logical namespace for all files. The result is an environment where file shares are available from any server node for any client. Virtualization use commodity operating system and hardware platform that offer modern networking capability (Figure 1)

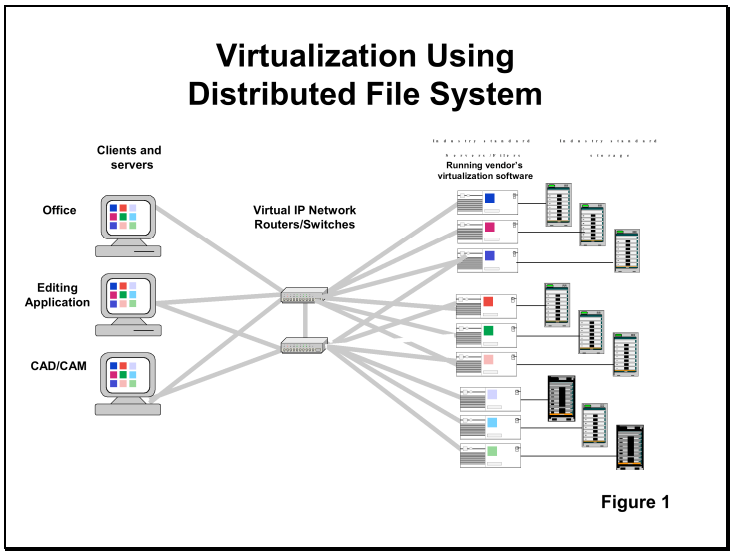

Figure 1

Products range from a "**software-only**" offer to a **complete system**. The "software-only" solution is mounted on a standard server provided by the client, and software agents are mounted on all servers that are aggregated. This solution enables the user to connect available server and storage resources and to select preferred storage vendor arrays as newer and more efficient configurations become available. The complete system solution contains a server or filer, the virtualization software and the storage arrays. Acquiring a complete system ensures that the vendor has performed complete qualification and will provide total support for all the components of the systems.

3- SELECTION CRITERIA

Virtualization provides great flexibility for dynamic management and allocation of devices and storage volumes that are not constrained by the physical dimensions of the devices. It allows users to grow or shrink the pool transparently as business needs require, and to migrate data within the pool dynamically to servers according to user policies. Hosts can be added or deleted. But products differ widely in their implementation and functionality

Five characteristics - availability, performance, manageability, scalability and cost – consistently rank at the top of the most wanted feature list. For a NAS virtualization implementation, two additional functionalities are commonly required – connectivity and ease-of-use. (Figure 2) The best way to measure the value of a product is to assess its features in the context of these characteristics.

- **Availability**

By far the most important user concern is full data and application availability. Beyond system and path redundancy and fail-over, data availability implies the ability to connect and disconnect on-line (hot plug) for non disruptive maintenance and reconfiguring, automatic fault detection, isolation and recovery, on-line reparability, and complete system restoration after fixing a failure. For a NAS virtualization solution, the following are a few questions that should be addressed:

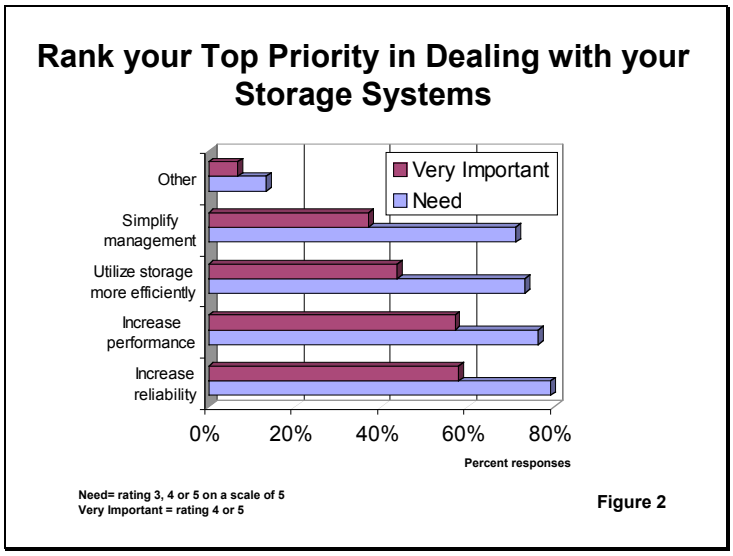

- Does the system present any single of point of failure
- Is a high availability cluster with automatic fail-over offered?
- Are disks/RAID arrays accessible and sharable between NAS heads
- Are hot spares used? How many?

- **Scalability**

Scalability is among the most important features required in a virtualization product (Figure 3) Scalability must allow growth in capacity, performance, connectivity and availability without affecting ongoing operations. Scalability includes adaptability and portability, for products and applications. Granularity of scalability may be important to your application.
Scalability NAS related questions include:

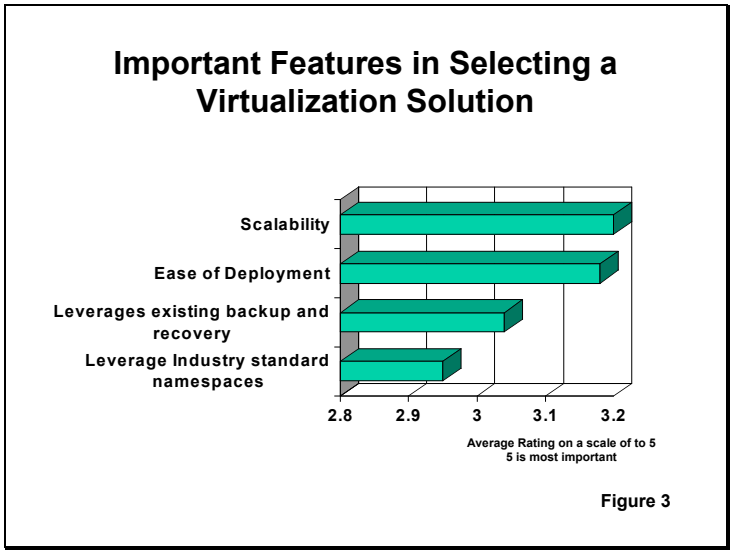

- Can the system independently scale capacity, performance, and I/O
- Can bandwidth or performance be increased without impacting clients?
- Does the system support heterogeneous arrays?
- Can a single file system be transparently scaled to multiple servers
- How many nodes and Ethernet ports are supported
- Does the system cover the longer term objectives for growth

- **Performance**

Improved performance is the top benfit users experience in virtualization (figure 4). Several levels of performance improvements can be expected, combining simultaneous data paths, real-time load balancing and dynamic path reallocation. Overall system performance depends on so many parameters , that there are few if any standards to equitably measure and compare performance. Main questions are:

- Is data migrated for capacity or performance / load balancing?
- What are the type, number and speed of processors used?
- What are the sizes of the write and read cache
- Is striping of files supported across server nodes?
- Is automatic load balancing performed?
- How many servers can be clustered?
- Does performance scale linearly with the number of nodes?

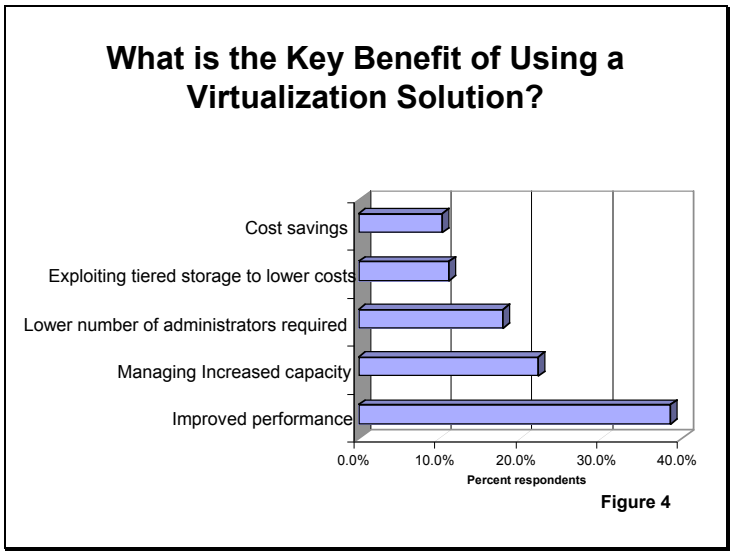

Figure 4

- **Manageability**

Managing storage is one of the most important costs in administering networks. The first important criterion in manageability is the ability to manage all storage components and servers from any point distributed on the network. Management automate the processes of configuration, monitoring, load balancing, diagnosing and reporting. The extent to which a NAS virtualization engine can integrate and automate some or all of these steps will determine its position in the hierarchy of product classes

- Is there consolidated management of many NAS?

- Is the management under Centralized or Distributed Control
- Can aggregation of anyone's NAS be performed, or is it limited to proprietary NAS
- Can open files be migrated transparently
- Does the system perform automatic discovery of new server and storage additions?
- Is a global namespace applicable?

- **Cost**

The real cost of storage is not in the hardware and software but primarily in the labor involved in managing storage and in the productivity loss. Therefore, the Total Cost of Ownership (TCO) needs to be taken into account, to include productivity gains due to increased performance, simplified management, better utilization of resources, the elimination of storage over-provisioning, and increased data availability. It has been estimated that with aggregation and automated management, the amount of storage manageable by a single administrator can be multiplied by a factor of 2 to 5

Questions to be addressed include:

- Is the solution based on low commodity hardware
- What is the Total Cost of ownership (TCO), how much of the operating cost can it save, and what is the Return On the Investment (ROI)
- How much additional hardware must be added to an Appliance or a software-only product, how long will the integration process take, and who will support it.
- Does the system include a comprehensive suite of offerings to simplify provisioning and maximize storage utilization
- Does the price include a portfolio of data management software for sharing, consolidating, and protecting data

- **Connectivity**

Connectivity features include the ability to add new elements without disrupting ongoing operations, added distance, serving a variety of platforms and operating systems. For an any-to-any connectivity, every host must be able to address every storage device on the network. If a switch is involved, parameters that affect connectivity include the number of ports and the ability for each port to connect to any other port in a duplex mode. NAS related questions include:

- Does the system allow connection to Windows, UNIX, Linux and Apple servers or workstations and can these operating systems seamlessly share files?
- Can it make use of available hardware and software resources
- What Disk Drive types are supported? (ATA, SATA, SCSI, FC, RAID)
- Can third party disk system be used?
- Can it aggregate other vendors' NAS systems
- Has Integration with leading backup, replication and/or ILM (Information Lifecycle Management) software been qualified?
- Can a cluster span geography? Is a WAN connectivity available?

Figure 5 rates the importance given by the users to these features

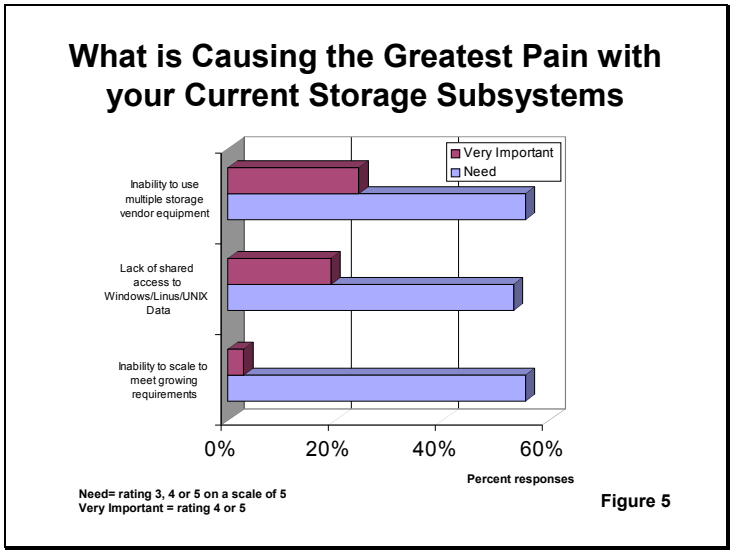

Figure 5

Associated Services

NAS virtualization products are often advertised with extended services options such as point-in-time data copying via snapshots, remote replication, mirroring, tiered storage management, and other utilities. Although storage virtualization facilitates these services, none are fundamentally dependent on storage virtualization technologies, and these services can easily be added at cost from reputable data protection software vendors. It is important nevertheless, for fair price comparison, to identify the extent to which these services are included.

4- CONCLUSION

NAS aggregation addresses the scaling, performance, and management problems that plague NAS scalability today. NAS virtualization has the potential to change the way users implement storage networks and share data. NAS file awareness provides a level of intelligence that is not achieved with SANs, supporting heterogeneous environments and solutions built on low-cost commodity hardware.

For the user, virtualization is fundamentally about simplification. Virtualization enables:

- Consolidated Management
- Sustained performance and availability for distributed storage
- Better utilization of servers, filers and storage resources
- Overall cost optimization.
- Improved data protection

Virtualization approaches differ greatly. Even when competing solutions appear to cover the same ground, subtle distinctions exist. One has to explore the actual architecture and the various features offered by the product in the context of the characteristics outlined above, for each user to prioritize his needs and optimize his choices.

Note: The figures are from a report titled "NAS and NAS virtualization" published by Peripheral Concepts, Inc. (www.periconcepts.com) and Coughlin Associates.