



Parallel NFS is the Future for Scale-Out Networked Storage



Networked storage allows administrators to manage “pool(s) of storage resources” providing a wide range of options that address operational costs, data protection and performance challenges.

By Don Grabski, Co-Chair, ESF NFS Special Interest Group.

What is the next practical networked storage industry standard that provides significant incremental value for customers and incentive for vendors to invest in networked storage solutions? The SNIA Ethernet Storage Forum NFS Special Interest Group believes that the Parallel Networked File System, or pNFS, could be just what is needed.

Scale out networked storage is a combination of key attributes of a distributed storage network topology that supports exponential growth in both performance and client connectivity with minimum administrative effort. Key attributes include the ability of the distributed network topology to easily expand and scale its compute, storage and file management resource pool without fork lift upgrades or increasing the management total cost of owner metric. The move toward the scale-out networked storage model represents a market evolution driven by field proven technology and a quantifiable return on investment for the customer. Scale-out networked storage represents the next logical phase of this market, because this architecture will provide a foundation to address the projected exponential data growth challenges across all market segments.

Challenges of Networked Storage Systems

Most customers have environments with independent networked storage systems that are not capable of achieving the ideal performance, capacity and client management utilization efficiencies they desire. Some of the challenges

include performance bottlenecks, non-scalable file systems, complex storage management and the fork lift upgrades required when installing new functionality or growing existing assets. Storage administrators are constantly looking for ways to reduce management costs, increase performance and lower total cost of ownership for networked storage purchases to address these challenges. Organizations on a whole want to reduce operational costs, increase productivity or solve a unique problem for a competitive advantage.

Network File System (NFS) addresses many of these challenges. NFS is a communication protocol that enables data stored on file servers available to any computer on a network. NFS clients are included in all common Operating Systems and allow servers to communicate with the file system in the storage. NFS also ensures interoperability among vendor solutions, allowing users to have a choice of best-of-breed products in their storage networks, and eliminating risks associated with proprietary technology. The availability of NFS v4.1 protocol (approved December, 2008) has resulted in many storage management enhancements, including global name space (a feature that can help storage administrators consolidate many pieces of hardware to look like a single system), head and storage scaling, as well as non-disruptive upgrades that can be made to the system while maintaining expected performance levels. These features all work to address the need to reduce storage operating costs,

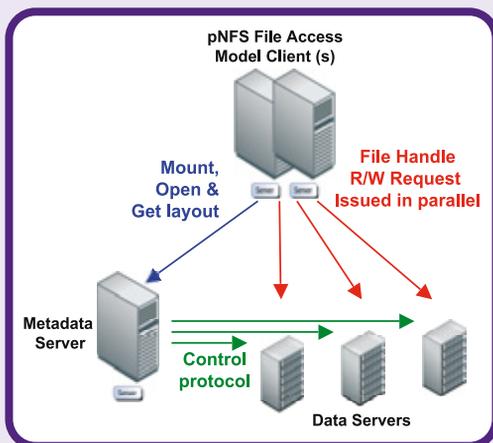
pNFS solves these issues by providing global name space functionality without requiring fork lift upgrades, all while allowing storage administrators to scale performance and storage capacity without disruption

improve storage network performance, and consolidate systems to reduce management overhead.

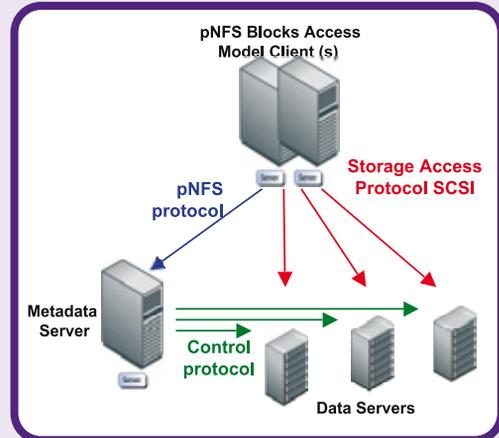
pNFS is the Future Protocol for Scale-Out Networked Storage

In order to remain viable and keep up with the requirements associated with scale-out networked storage, the current NFS standard protocol must make some dramatic changes. The future protocol for scale-out networked storage is Parallel I/O Network File System (pNFS). pNFS allows users to access storage devices directly, and in parallel, by leveraging the combination of Parallel I/O and NFS. It supports block, file and object-based data files. Parallel I/O delivers very high application performance and allows for massive scalability without diminished performance. Single sequential I/O patterns create bottlenecks which adversely affect performance — for example load balancing is not available and it can't aggregate other devices. However, pNFS solves these issues by providing global name space functionality without requiring fork lift upgrades, all while allowing storage administrators to scale performance and storage capacity without disruption. This lowers the total cost of ownership and maximizes consolidation of storage. pNFSv4.1 with pNFS support is expected to be available in 2010.

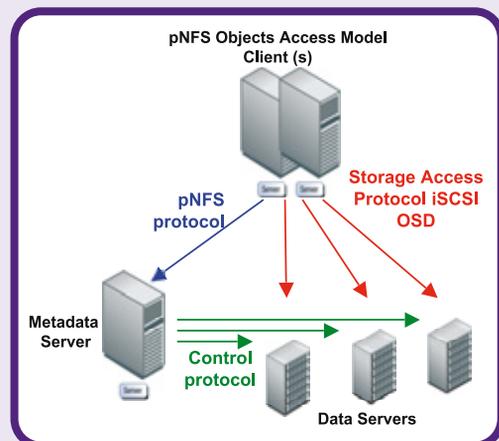
How pNFS Works:



- Client mounts and opens a file on the server
- Servers grant the open and a file stripe map (layout) to the client
- The client can read/write in parallel directly to the NFSv4.1 data servers



- Client mounts and opens a file on the server
- Servers grant the open and a block map (layout) to the client
- Based on the layout obtained (read or write), the client can read/write in parallel directly to the SCSI target's



- Client mounts and opens object
- Servers grant the open and an object stripe map with object capabilities (layout) to the client
- Based on the layout obtained (read or write), the client can read/write in parallel directly to the OSD targets

For more information regarding SNIA NFS Special Interest group, sponsored educational events and resources visit http://www.snia.org/forums/esf/programs/about/nfs_sig