



pNFS Status


September 2010

Larry Jones, Panasas

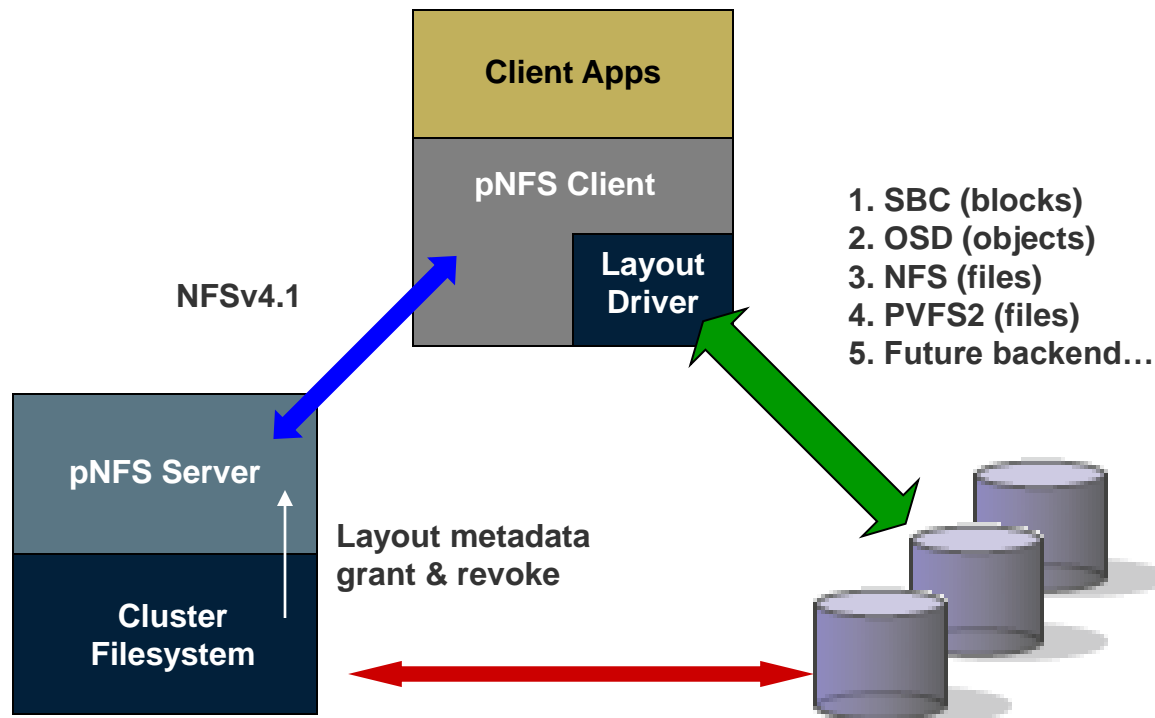


Review: The pNFS Standard

- The **pNFS** standard defines the NFSv4.1 protocol extensions between the **server and client**
- The **I/O** protocol between the **client and storage** is specified elsewhere, for example:
 - SCSI **Block** Commands (**SBC**) over Fibre Channel (**FC**)
 - SCSI **Object**-based Storage Device (**OSD**) over iSCSI
 - Network **File** System (**NFS**)
- The **control** protocol between the **server and storage** devices is also specified elsewhere, for example:
 - SCSI **Object**-based Storage Device (**OSD**) over iSCSI



- Common client for different storage servers



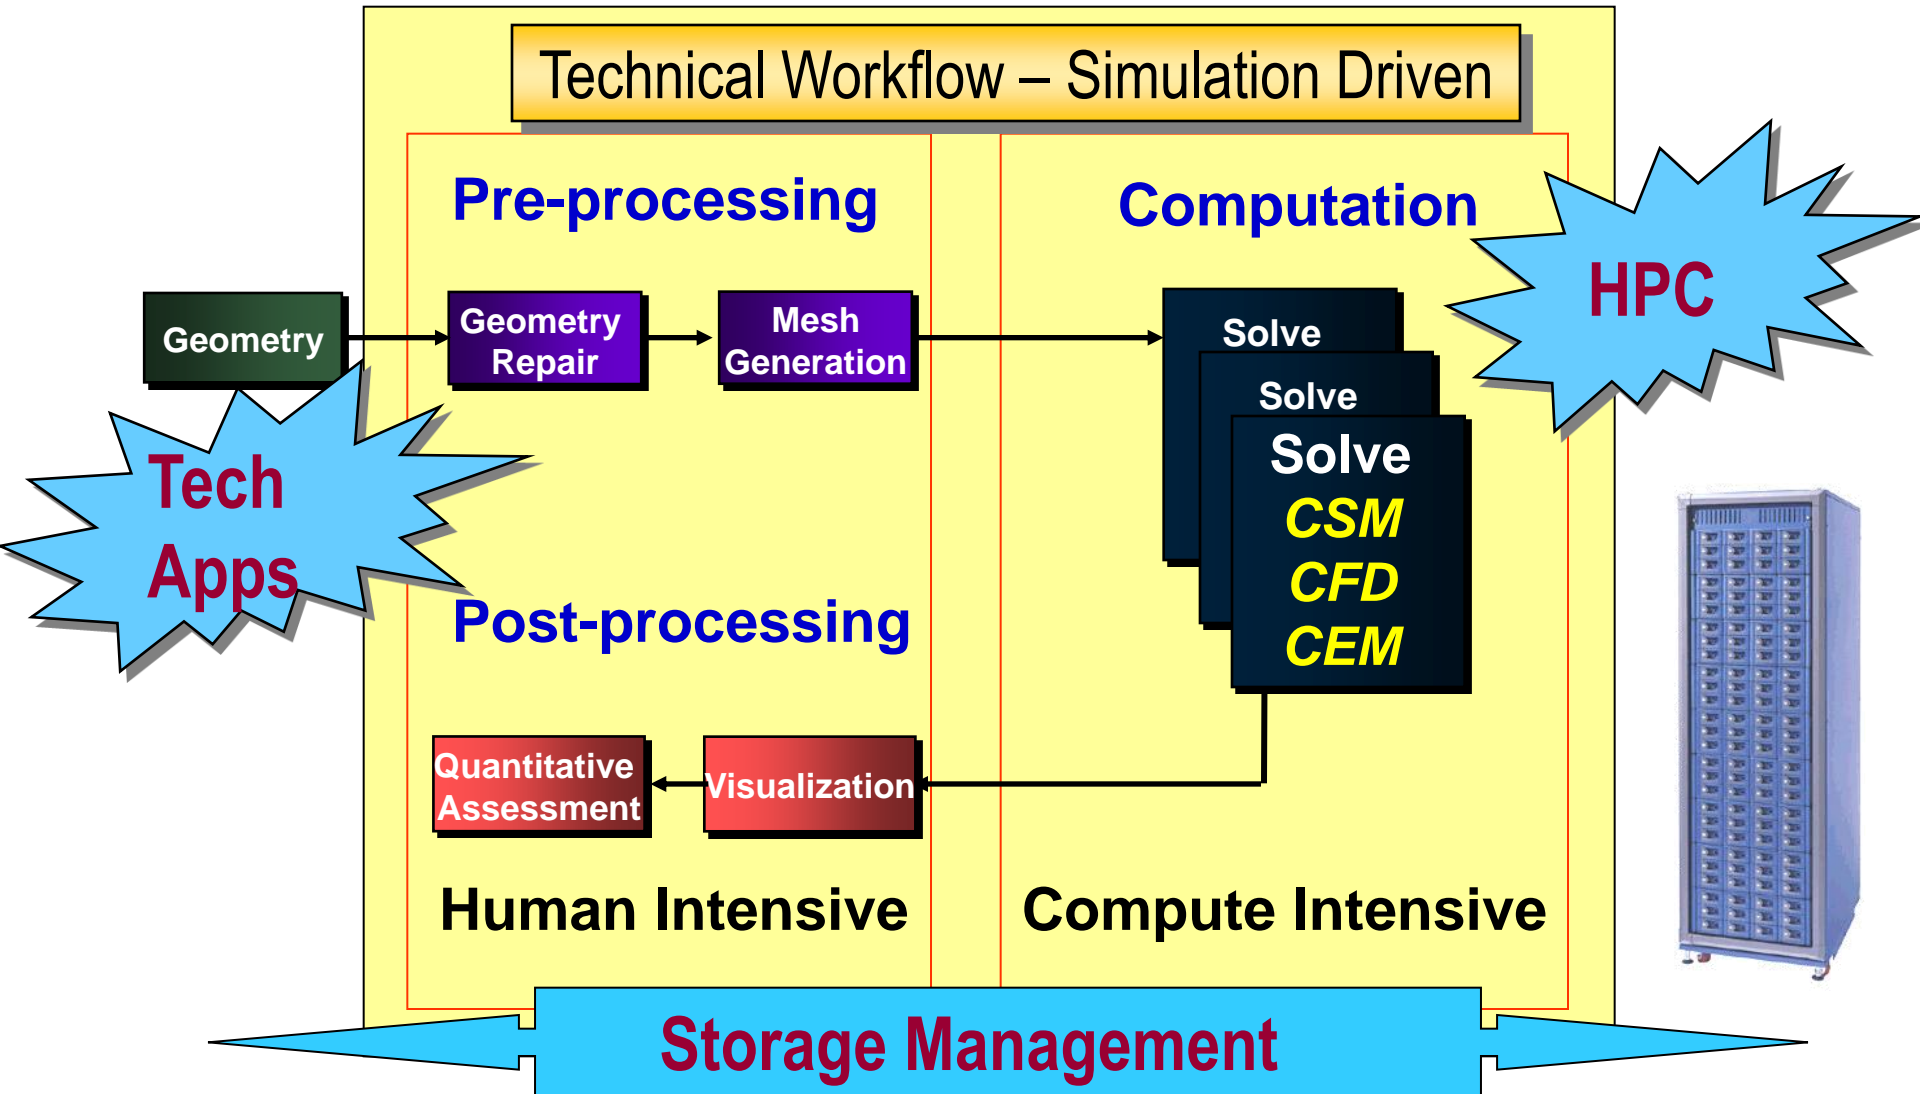
Standards drive adoption and unlock markets



Standards drive down costs



pNFS adds value across the workflow



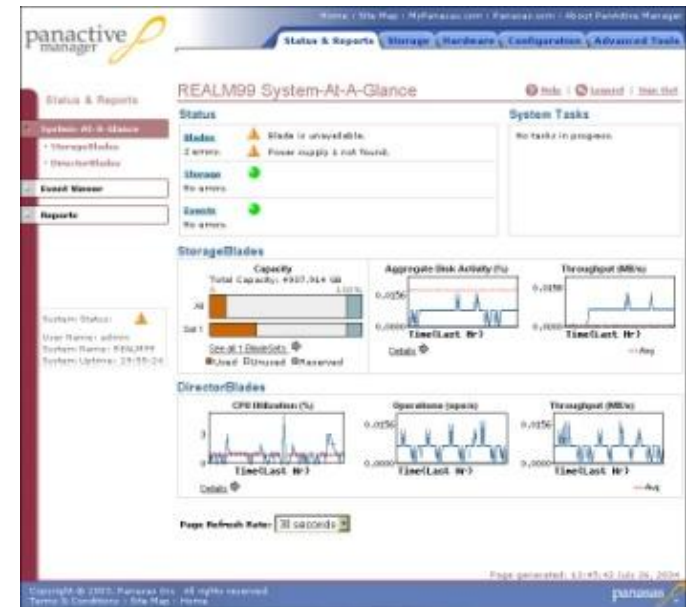
- Scalability
 - Consolidate more users into a single system
 - Accelerate workflows by consolidating data across the organization
- Ease of Use
 - High performance, standard pNFS client delivered with the desktop OS
 - Single mount point makes data sharing easy across projects and workgroups
- Performance
 - Faster I/O to high performance workstations – eliminate the “filer bottleneck”
 - Load balance small I/O requests across all servers

Performance (YMMV)	NFS	pNFS
High Concurrency (many clients/files)	10s to 100s of users	Scales linearly
Single Client Read/Write	10s to 100s of MB/s	Scales linearly
Capacity	10s to 100s of TBs	Scales linearly

- Scalability
 - Eliminates islands of storage - NFS
 - Eliminates store and forward protocols – Clustered NFS
- Performance
 - Enables concurrent write to eliminate serial I/O write bottlenecks
 - Eliminates hot spots due to poor data distribution across servers
 - Balance I/O demand across many storage devices
 - Scale I/O to meet the most demanding cluster requirements

Performance (YMMV)	NFS	pNFS
Concurrent Write (N ->1)	Difficult	Scales linearly
Concurrent Read	10s to 100s of MB/s	Scales linearly
High Bandwidth – Read/Write	10s to 100s of MB/s	Scales linearly

- Standard client
 - Transparent to applications
 - Common client for different storage back ends
 - Fewer support issues for storage vendors
 - Normalizes access to clustered file systems
- Manageability
 - No Mount load balancing
 - No storage load or capacity balancing
 - No (fewer) Mount Tables to manage
 - Easier provisioning – Single storage pool
- Scalability
 - Eliminates store and forward protocols
 - Virtualized storage - distributes data across storage cluster



- 2003 First pNFS meeting among vendors
- 2005 First IETF drafts
- 2008 Approval of drafts for standard track
- 2010 RFC status achieved!
 - 5661: NFSv4.1 protocol
 - 5662: NFSv4.1 XDR Representation
 - 5663: pNFS Block/Volume Layout
 - 5664: pNFS Objects Operation

- 2.6.34
 - Merge window February 2010, Released May 2010
 - 21 NFS 4.1 patches
- 2.6.35
 - Merge window May 2010, release August 2010
 - 1 client and 1 server patch (4.1 support)
- 2.6.36
 - Merge window August 2010
 - 16 patches accepted into the merge
- 2.6.37 preparations
 - 290 patches represent pNFS functionality
 - Finalizing patches before October Bake-a-thon testing session

- 2.6.37
 - Merge window November 2010
 - Files pNFS client and server
- 2.6.38
 - Merge window February 2011
 - Object pNFS client and server
- 2.6.39
 - Merge window May 2011
 - Blocks client and server

Key pNFS Participants



- Panasas (Objects)
- Network Appliance (Files over NFSv4)
- IBM (Files, based on GPFS)
- EMC (Blocks, HighRoad MPFSi)
- Sun/Oracle (Files over NFSv4)
- U of Michigan/CITI (Files over PVFS2)



- RedHat:

- Red Hat is fully committed to support pNFS in an upcoming RHEL6.x release.
- Red Hat is actively contributing to get pNFS working in the upstream kernel.
- Red Hat has built a special version of pNFS for Fedora and is working to test that with multiple partners.
- Sayan Saha – RedHat product manager – 09/10/10



- Microsoft / CITI

- CITI would like to announce that we have a beta version of the pnfs windows client available for download. You can find the location of the git repository as well as directions on how to build and install it at <http://citi.umich.edu/projects/nfsv4/windows/>.
- --CITI – 9/13/10



- Panasas

- Panasas was instrumental in creating the standard, has developers working on the kernel implementation of pNFS and will deliver pNFS servers based on our industry-leading PanFS products when the Linux clients are available.
- Faye Pairman, CEO – 08/10/10



- NetApp

- "The pNFS protocol will also be a boon to server clustering in general. "You get all the management advantages of storage clustering without having to pay for any of the performance drawbacks that you would get in the absence of a more agile protocol like parallel NFS."
- Mike Eisler, Sr. Technical Director – 06/07/10

- EMC

- From the Celerra VG8 spec sheet:
- DART file server facilities
- Protocols supported:
 - NFSv2, v3 and v4 (including NFSv4.1 with pNFS support)



- BlueArc

- “the scale aspect of BlueArc products are going to be enhanced with Parallel NFS (pNFS) support, with BlueArc likely to introduce a PAFS (Parallel Aggregate File System) capability to its software.”
- Mike Gustafson, CEO - The Register – 3/24/10



Thank you for supporting pNFS!

- pNFS is a long range investment for the Technical Computing and storage community
- SNIA is sponsoring a Developer/Future End User Interest Group
 - Looking for potential pNFS Users and Analysts that will provide feedback to the vendor community and/or talk to the press
 - See me if you're interested...

