NFSv4 and Data Management

Andy Adamson
CITI, University of Michigan
andros@umich.edu

NFSv4 – Not Your Father's NFS

- Strong, mandatory security
- NFSv4 supports ACLs
- File System redirection for Global Naming, migration, and replication
- Extensible by design
- Other cool features

NFSv4 – Strong Security

- RPCSEC_GSS
  - ONC/RPC + GSS API that supports data privacy and integrity
- Mandatory GSS security mechanisms
  - Kerberos v5
  - SPKM-3 (X.509 based)
  - LIPKEY (think TLS)

NFSv4 – Flexible Trust Virtualization

- NFSv4 ACLs
  - Modeled on Windows ACLs, superset of POSIX ACLs
  - Like Globus GSI, NFSv4 puts names not numbers on the wire
  - File system ACL identity - UID/GID
  - NFSv4 maps names to UID/GID
- NFSv4 name-ID mapping enables foreign user and group support
  - Foreign users/groups in ACLs
  - Grid DN and groups can be used.
- Mapping service can be managed by network, storage, or collection administrators.

NFSv4 and Data Virtualization

- File System redirection separates name space from storage
  - Client queries server for location(s) of file system at any time or in response to an NFS4ERR_MOVED
  - Sever returns location of replicated or migrated file systems.
- Replication and migration features expanded in NFSv4.1
  - New fs_locations_info proposed
- NFSv4 global file name space
  - Location information enables the construction of a global or organization wide file name space
- Server to Server Migration – Replication protocol is under consideration

pNFS – NFSv4.1 Feature

- NFSv4.1 pNFS removes single client to single server bottleneck
- Divides the NFSv4 protocol into control and data paths
- Direct I/O to storage
  - file, object, and block I/O
- Stripe READ/WRITE across multiple servers
- Allows cluster file systems to join enterprise NFSv4 name space
- Linux and Solaris pNFS client prototypes
- GPFS, Network Appliance, Panasas, EMC, Sun, Lustre are all implementing pNFS server prototypes

pNFS Architecture

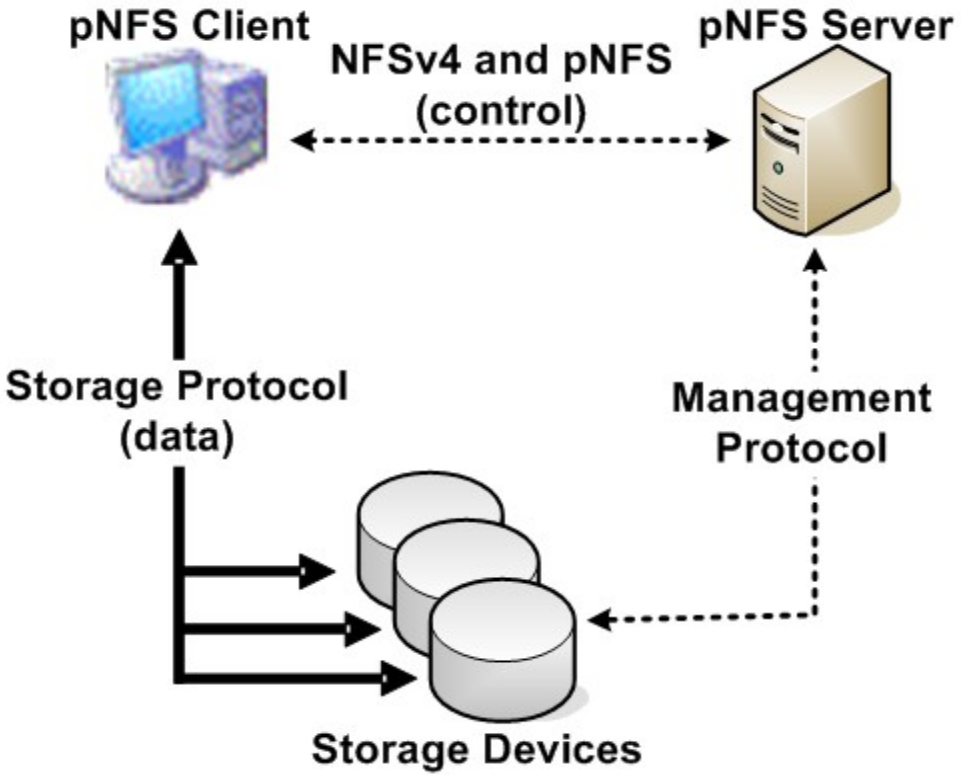

NFSv4 and Data Management

- NFSv4 provides these features to mange geographically separated storage repositories:
  - Separation of namespace from storage: replication/migration/global name space
  - On the wire security with data integrity, privacy
  - Trust virturalizaion via name-ID mapping
- NFSv4 can be an important part of Data virtualizaion.
- NFSv4 can glue together Storage virtualization solutions.