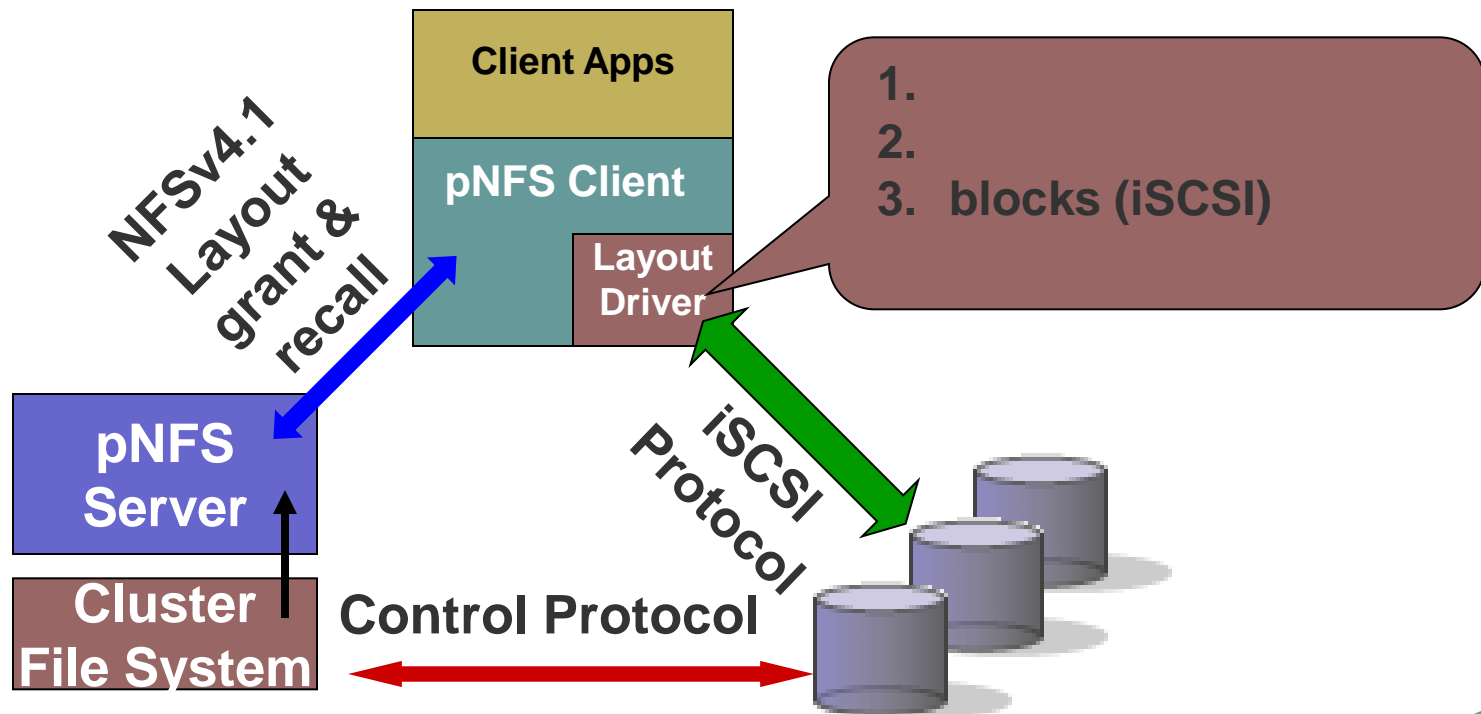


# pNFS Block Status - EMC

# Linux pNFS Block Layout

- Transparent to applications
- Common client for different storage vendors
- Fewer support issues for storage vendors
- Normalizes access to clustered file systems



# NFSv4.1 – with pNFS Block

- Linux pNFS block layout Client
  - Support generic layout – CITI Univ. of Michigan
  - Support block layout client (iSCSI) - F-16 CITI
  - Maintenance of block layout – CITI via Linux kernel Bugzilla
  - Performance monitoring and patching – CITI via Bugzilla
  - New EMC VNX release with improved pNFS block performance – Q3 2012
    - High bandwidth
    - Improved small files performance
- Support to RedHat
  - Support block layout – EMC Elab work with RedHat to qualify pNFS
  - EMC Elab will qualify Fedora 16

# pNFS Block – Protocol enhancement

- NFSv4.1/pNFS were standardized at IETF
  - NFSv4 working group (WG)
- New Disk Protection draft approved by IESG
  - Extension to RFC 5663
  - Added mechanism that enables identification of block storage devices used by pNFS to be protected using GPT
- EMC VNX will implement it when RFC is approved

# pNFS Block Performance

- Latest Performance Using 10GbE server iSCSI and 1 GbE clients 4 threads per client
  - Single Client
    - Write: 107 MB/sec
    - Read: 90 MB/sec
  - 8 Clients
    - Write: 852 MB/sec – 98% scalability
    - Read: 674 MB/sec – 96% scalability