



pNFS in DPM

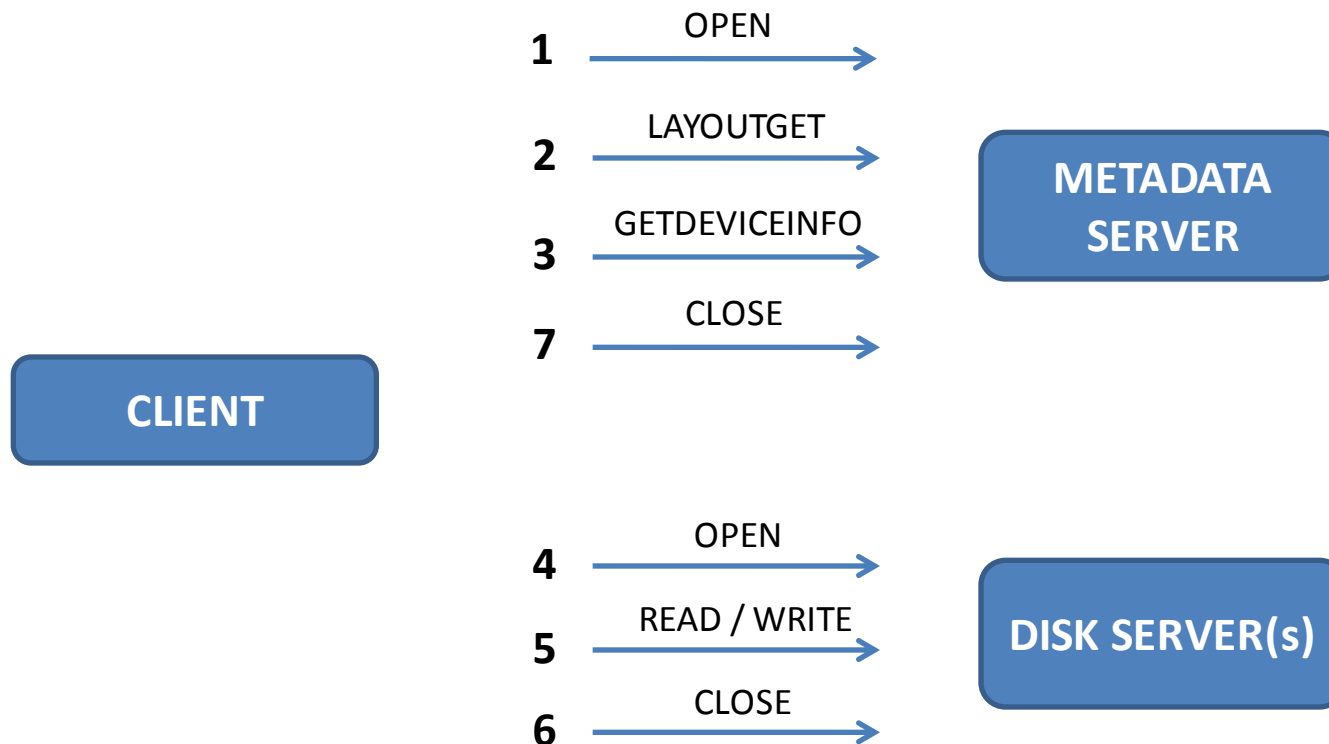
Ricardo Rocha

(on behalf of the DPM team)

NFS 4.1/pNFS: Why?

- Industry standard (IBM, NetApp, EMC, ...)
- No vendor lock-in
- Free clients (with free caching)
- Strong security (GSSAPI)
- Parallel data access
- Easier maintenance
- ...
- But you know all this by now...

NFS 4.1/pNFS: Overview



<https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/NFS41>

pNFS in DPM

Status & Performance

Client Availability


- pNFS support in linux kernel from $\geq 2.6.38$
- nfs-utils $\geq 1.2.3$
- Latest Fedora and Debian Sid have these
- But the ones we support do not
 - RHEL6: 2.6.32
 - Debian 6/Squeeze: 2.6.32
- We provide packages for RHEL5 / RHEL6
 - Enabled pNFS in the elrepo mainline kernel
 - nfs-utils and AFS module we package ourselves

Implementation Status

- Fully functional
- Deployed at a few sites as a beta
 - ASGC, UK GridPP, ...
- New beta release every month
 - Announced at the DPM users forum
 - Already fixed most issues seen regarding metadata operations, performance
- This will evolve into a production state soon
 - read-only for a start?

First Performance Results


IOZONE Results



- Server
 - Xeon 4 Cores 2.27GHz
 - 12 GB RAM
 - 1 Gbit/s links
- Client
 - Dual core
 - 2 GB RAM
 - 100 Mbit/s link

First Performance Results

NFS vs RFIO



- Server
 - Xeon 4 Cores 2.27GHz
 - 12 GB RAM
 - 1 Gbit/s links
- Client
 - Dual core
 - 2 GB RAM
 - 100 Mbit/s link
- 8 KB block sizes

Issues & Roadmap

- First we need...
 - Tests with a faster network link
 - Testing with a larger number of concurrent clients
 - WAN testing
 - Enable bigger block sizes
- We're using a test cluster in Taipei for the large scale tests (1000 cores)

Issues & Roadmap

- And then...
- X509 certificate support
 - Still not figured out... needs a strong focus
 - What we need is a X509 GSSAPI plugin, integrating with linux gssmech library
 - Showstopper?
- Kernel availability?
- Further validation with other implementations