

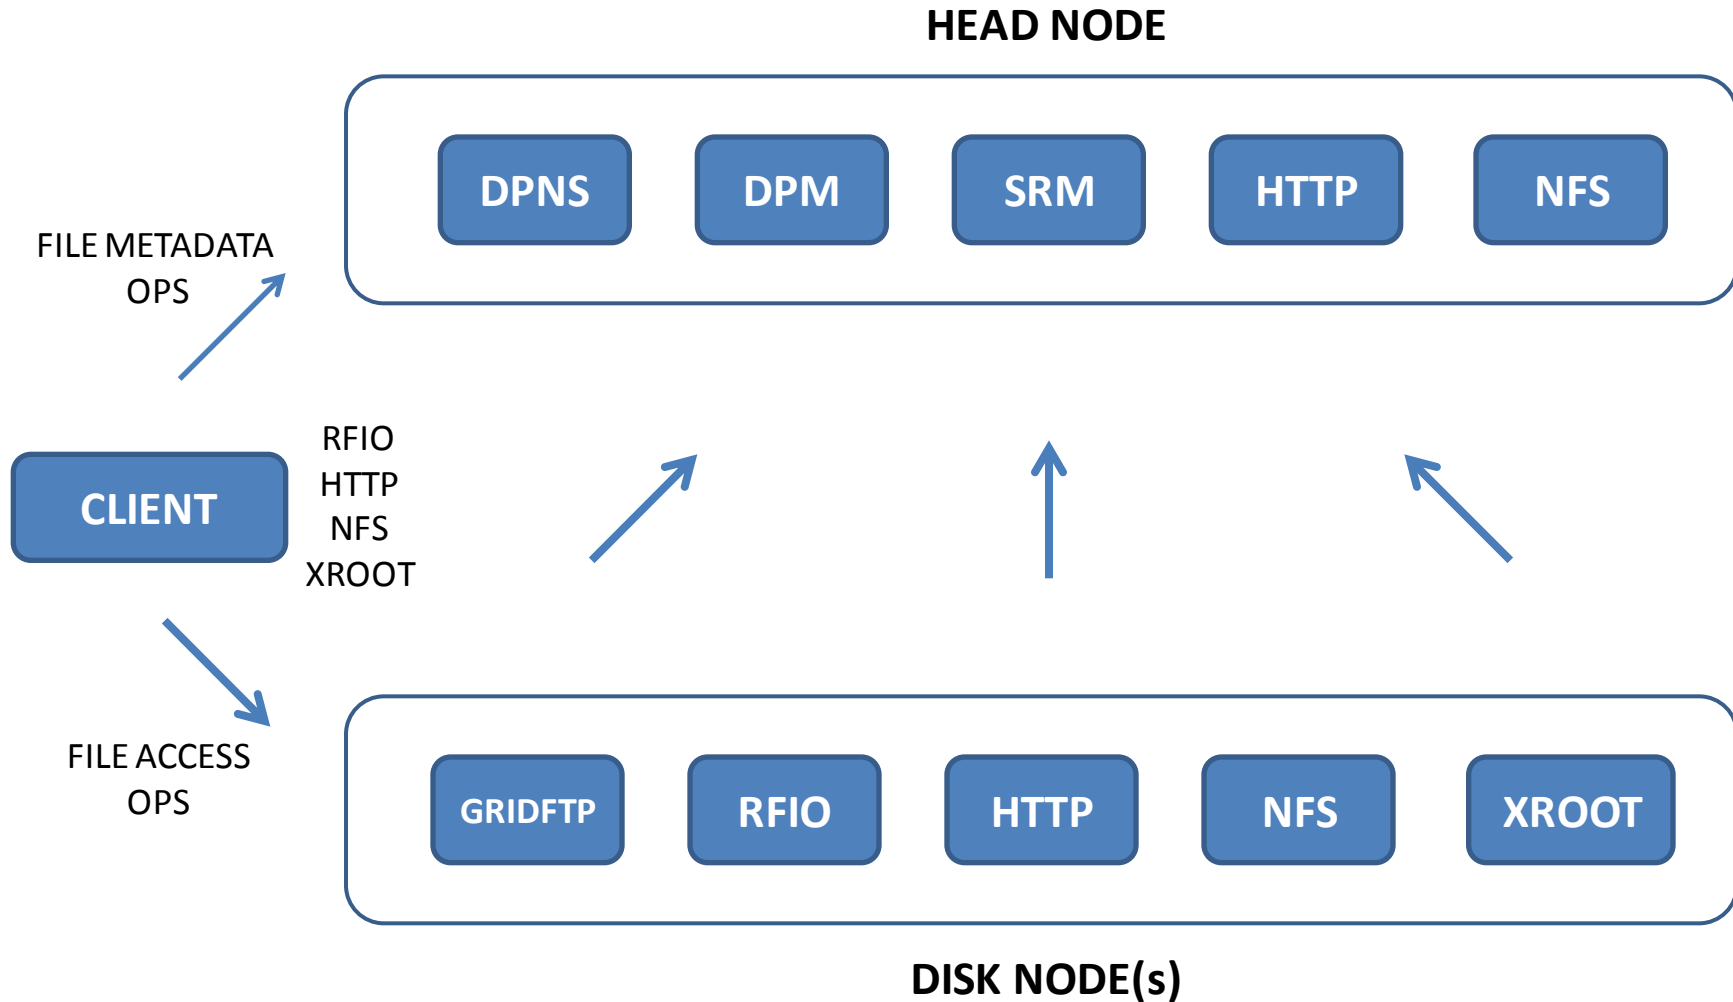


# Standards, Status and Plans

Ricardo Rocha

( on behalf of the DPM team )

# DPM Overview



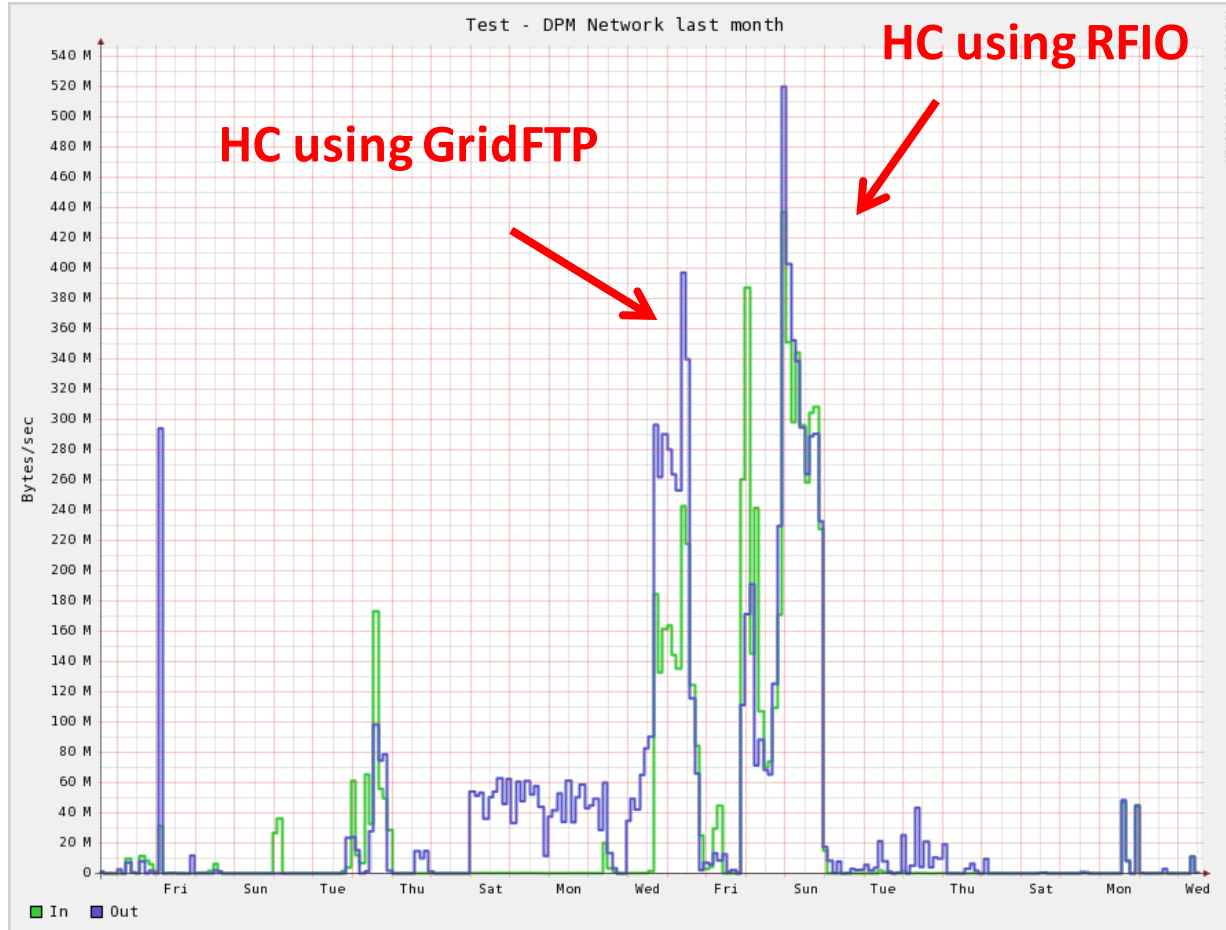
# DPM Core

## 1.8.2, Testing, Roadmap

## DPM 1.8.2 – Highlights

- Improved scalability of all frontend daemons
  - Especially with many concurrent clients
  - By having a configurable number of threads
    - Fast/Slow in case of the dpm daemon
- Faster DPM drain
  - Disk server retirement, replacement, ...
- Better balancing of data among disk nodes
  - By assigning different weights to each filesystem
- Log to syslog
- GLUE2 support

# DPM Core : Extended Testing Activity



Cluster at ASGC  
(thanks!)

1000 Cores

Regular  
Hammercloud  
Tests

*Thanks to ShuTing for the plots ( preliminary results )*

## DPM Core – Roadmap

- Package consolidation: EPEL compliance
  - Fixes in multi-threaded clients
  - Replace httpg with https on the SRM
  - Improve dpm-replicate (dirs and FSs)
  - GUIDs in DPM
  - Synchronous GET requests
  - Reports on usage information
  - Quotas
  - Accounting metrics
  - HOT file replication
- 1.8.3
- 1.8.4
- 1.8.5

# DPM Beta Components

HTTP/DAV, NFS, Nagios,  
Puppet, Perfsuite, Catalog  
Sync, Contrib Tools

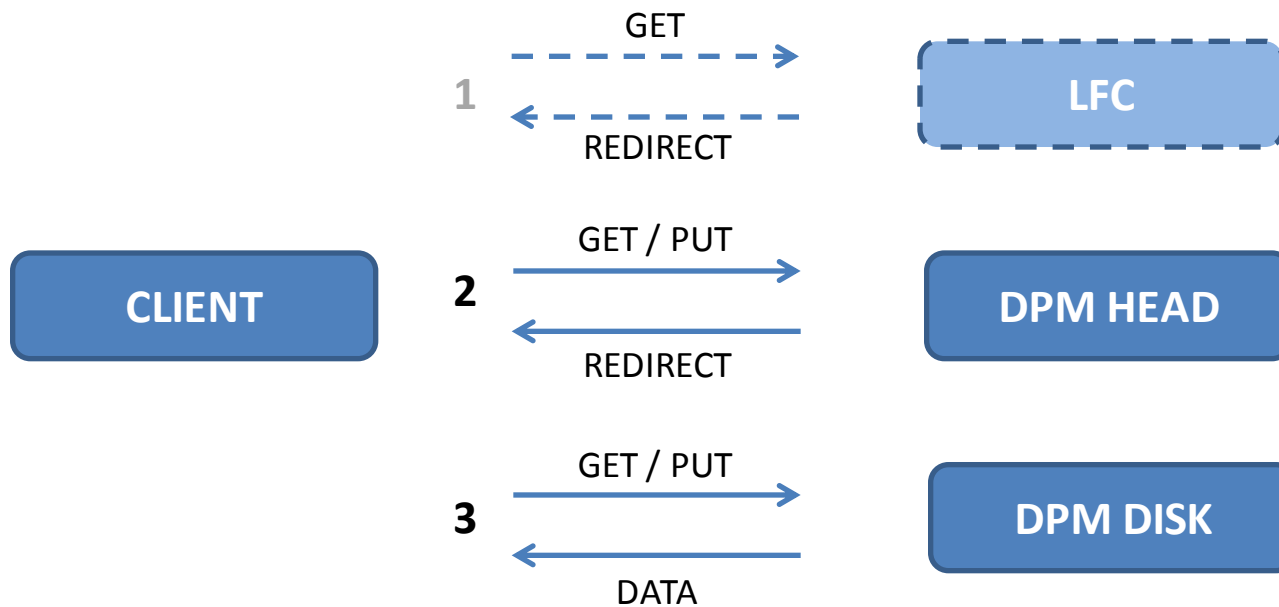
<https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Dev/Components>

# DPM Beta: HTTP / DAV

## Overview, Performance



# DPM HTTP / DAV: Overview



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

# DPM HTTP / DAV: Overview



The screenshot shows a web browser window with the address bar containing `https://belit.cern.ch/dpm/cern.ch/home/dteam/`. The main content area displays a directory listing for `/dpm/cern.ch/home/dteam/` with the following entries:

Permissions	UID	GID	Size	Timestamp	File Name
-rw-rw-r--	101	101	28.8K	2011-07-08 T 13:16:13	<a href="#">WebDAV.odp</a>
drwxrwxr-x	0	101	63	2011-07-08 T 10:46:45	<a href="#">generated/</a>
-rw-rw-r--	101	101	175.0K	2011-07-08 T 12:32:13	<a href="#">image.jpg</a>

Below the listing, the following text is displayed:

*Request by /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=aalvarez/CN=678984/CN=Alejandro Alvarez Ayllon (No proxy)*  
*Powered by DPM-WebDAV*

## HTTP : Client Support

	curl	browser
OS	Any	Any
GUI	NO	YES
CLI	YES	NO
X509	YES	YES
Proxies	YES	Only IE so far
Redirect	YES	YES
PUT	YES	NO

- **Recommendation:** browser/curl for GET, curl for PUT
- Chrome Issue 9056 submitted for proxy support

## DAV : Client Support

	TrailMix	Cadaver	Davlib	Shared Folder	DavFS2	Nautilus	Dolphin
OS	Firefox < 4	*nix	Mac OS X	Windows	*nix	Gnome	KDE
GUI	YES	NO	YES	YES	N/A	YES	YES
CLI	NO	YES	NO	NO	N/A	NO	NO
X509	YES	YES	NO	YES	YES	NO	NO
Proxies	?	NO	NO	YES	NO	NO	NO
Redirect	YES	NO	YES	Not PUT	NO	NO	YES

- Updated analysis based on initial one from dCache
- **Recommendation:** Cadaver for \*nix, Windows explorer

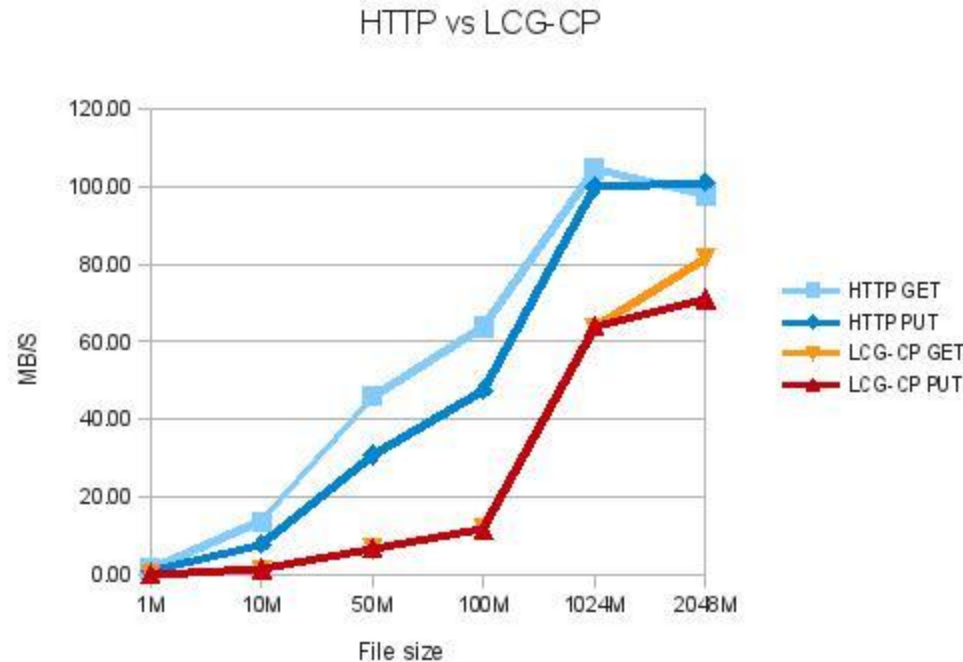
## HTTP vs GridFTP : Multiple streams

- Not explicit in the HTTP protocol
- But needed for even higher performance
  - Especially in the WAN
- So we added it, with some semantics
  - Small wrapper around libcurl
  - PUT with '0 bytes' && null content-range == end of write
- Submitted patch to libcurl to allow ssl session reuse among parallel requests

## HTTP vs GridFTP: 3<sup>rd</sup> Party Copies

- Implemented using WEBDAV COPY
- Requires proxy certificate delegation
  - Using gridsite delegation, with a small wrapper client
- Requires some common semantics to copy between SEs (to be agreed)
  - Common delegation portType location and port
  - No prefix in the URL ( just `http://<server>/<sfm>` )

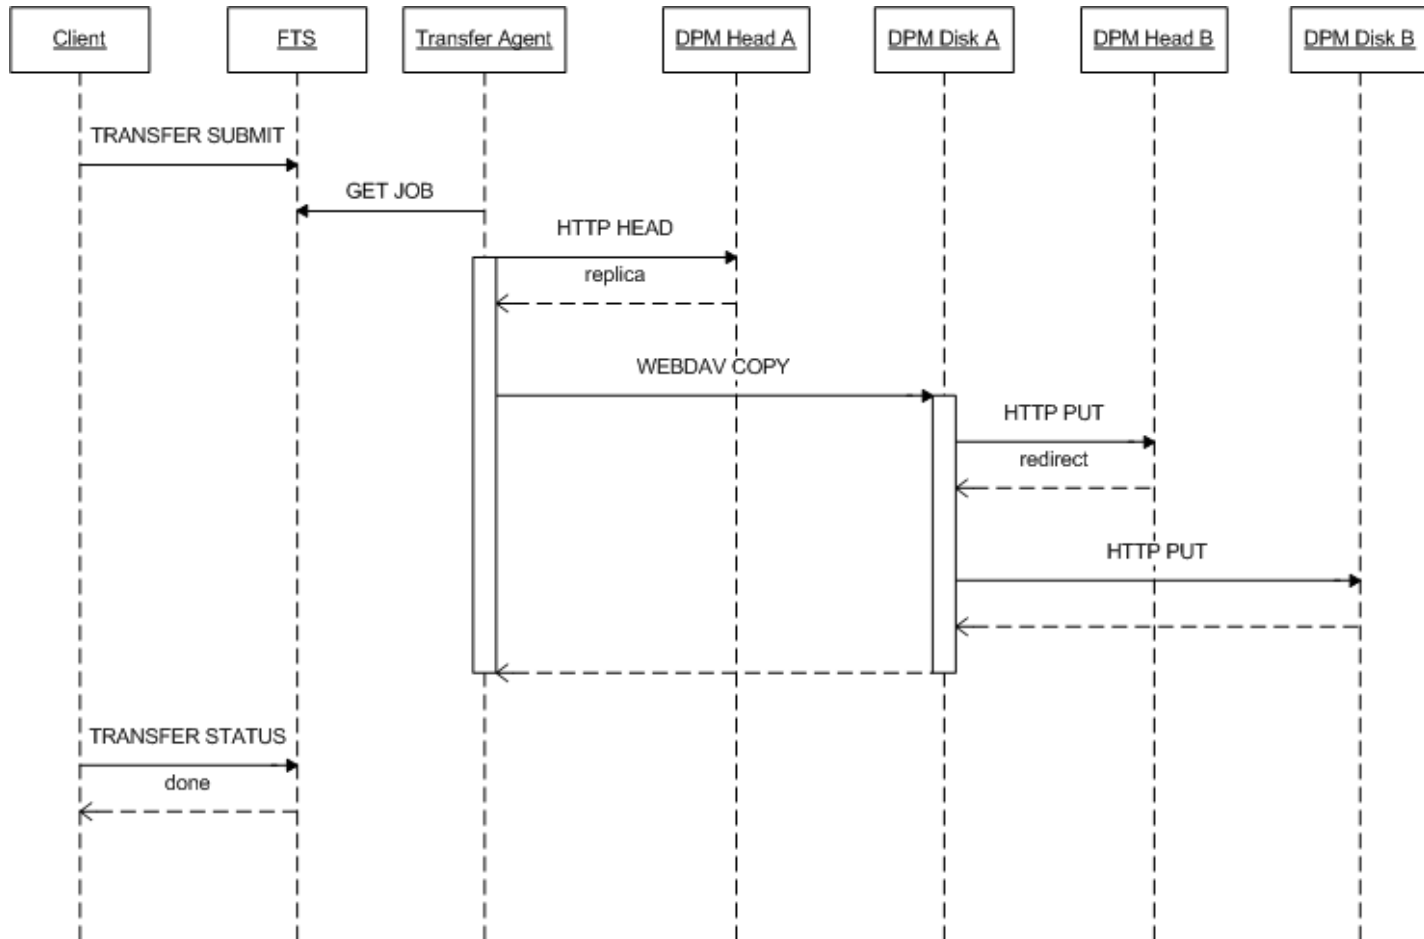
## DPM HTTP / DAV : Performance



- Xeon 4 Cores 2.27GHz
- 12 GB RAM
- 1 Gbit/s links

- No difference detected in LAN with different number of streams
  - But early results do show a big difference on the WAN
- lcg-cp configured to use gridftp
- File registration & transfer times considered in both cases

# DPM HTTP / DAV : FTS Usage



Example of FTS usage



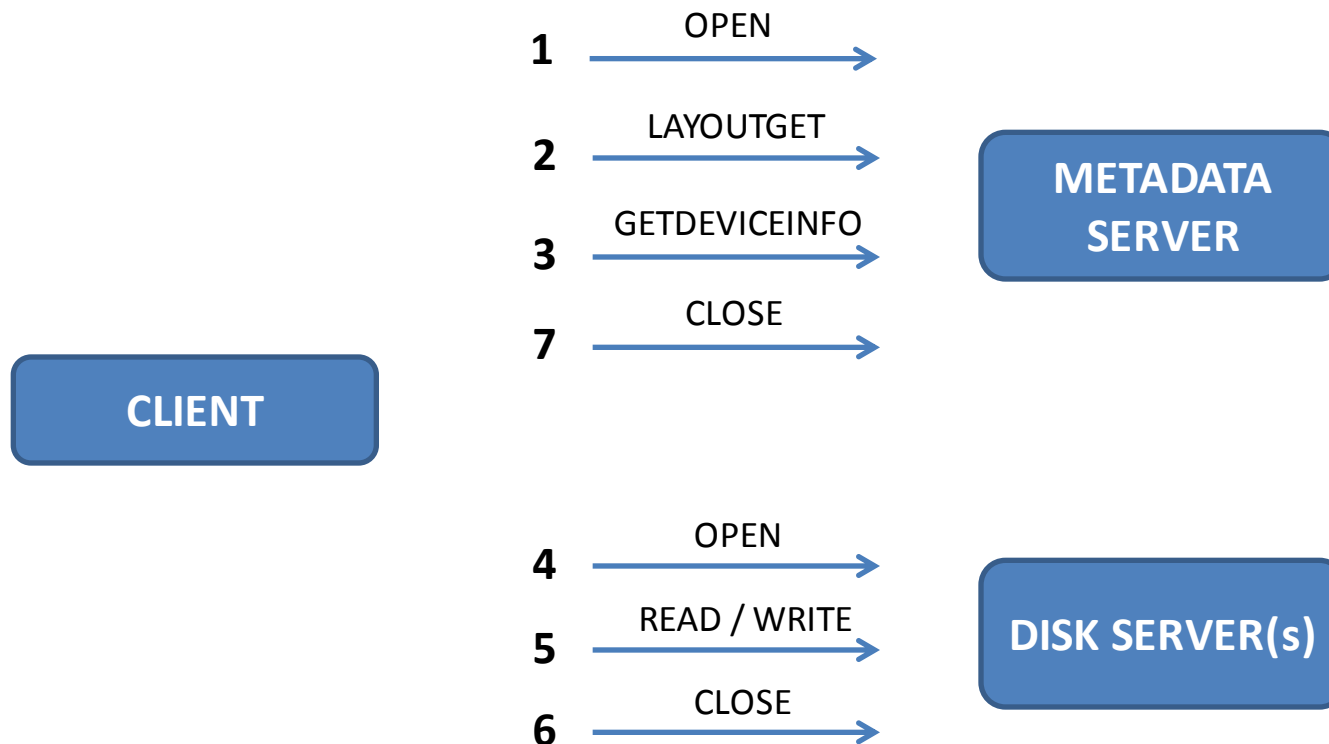
# DPM Beta: NFS 4.1 / pNFS

## Overview, Performance

## NFS 4.1/pNFS: Why?

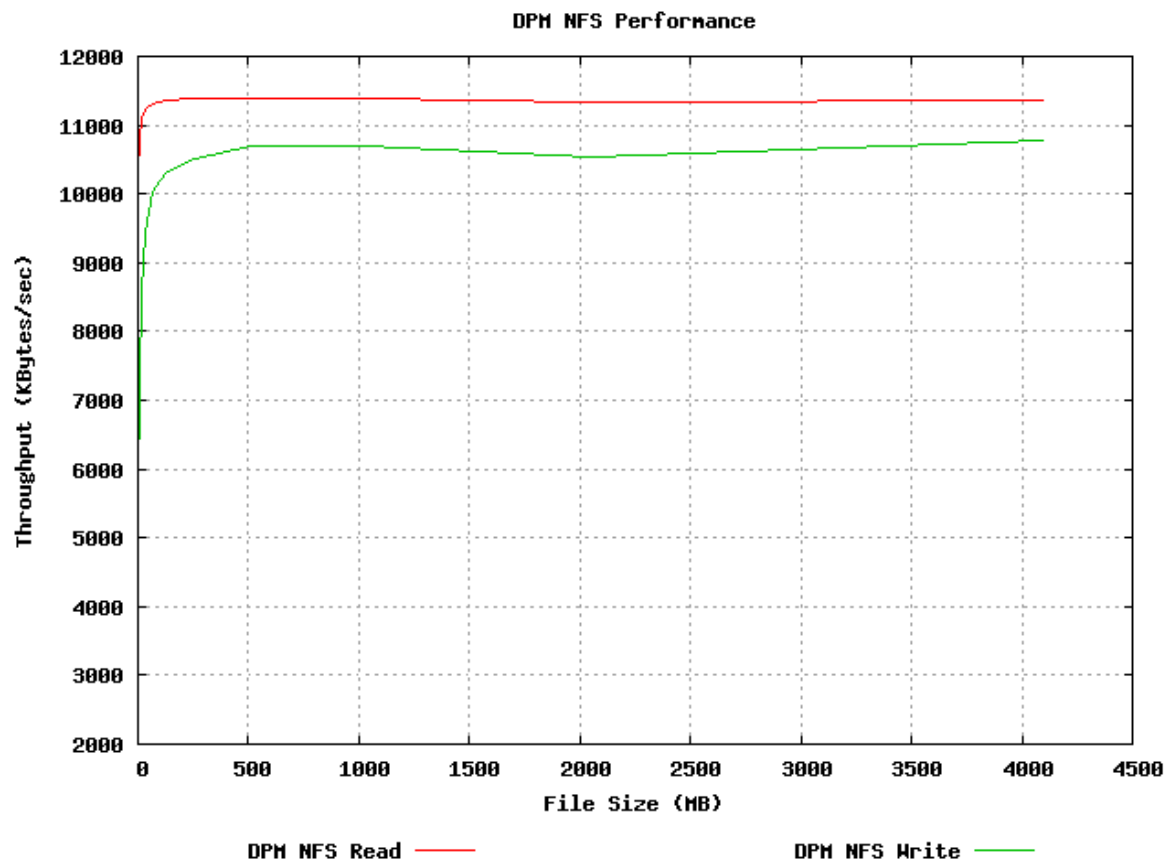
- Industry standard (IBM, NetApp, EMC, ...)
- No vendor lock-in
- Free clients (with free caching)
- Strong security (GSSAPI)
- Parallel data access
- Easier maintenance
- ...
- But you know all this by now...

# NFS 4.1/pNFS: Overview



# NFS4.1 / pNFS: Performance

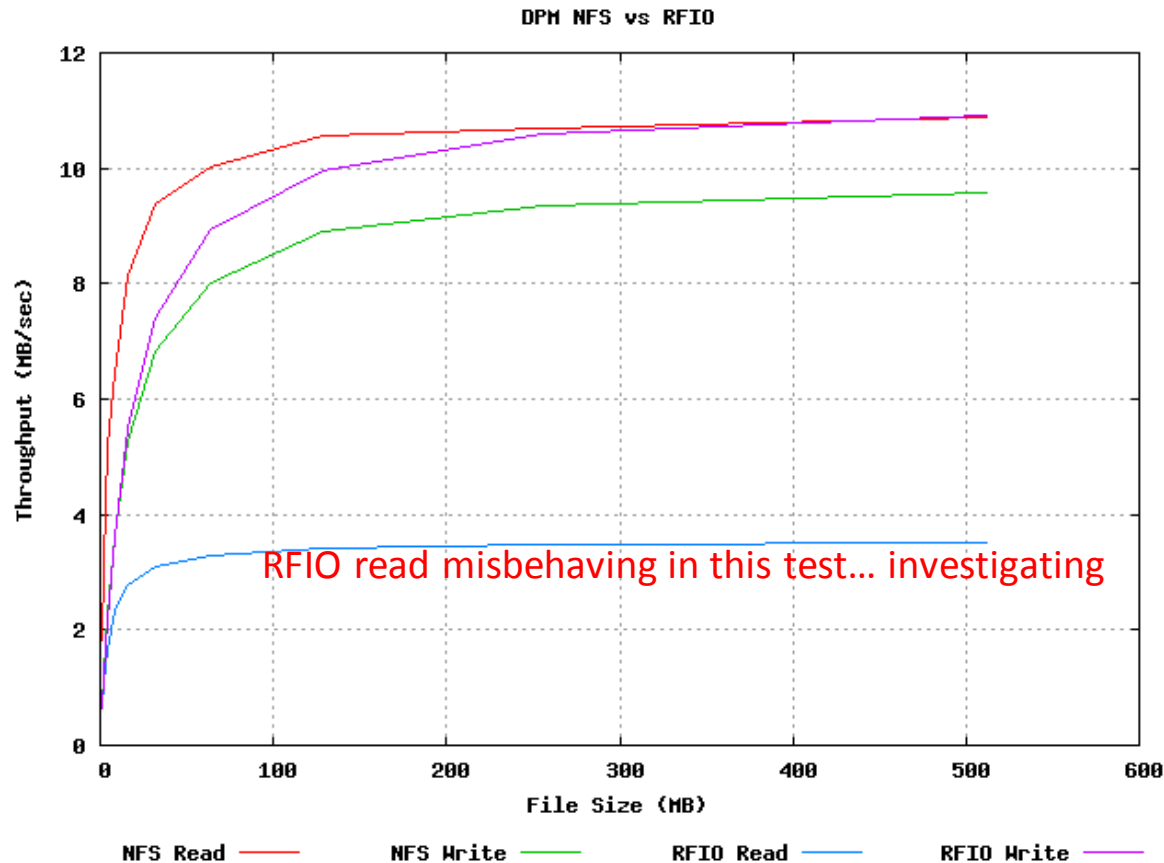
## IOZONE Results



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

# NFS4.1 / pNFS: Performance

## 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

## Conclusion

- 1.8.2 fixes many scalability and performance issues
  - But we continue testing and improving
- Popular requests coming in next versions
  - Accounting, quotas, easier replication
- Beta components getting to production state
  - Standards compliant data access
  - Simplified setup, configuration, maintenance
  - Metadata consistency and synchronization
- And much more extensive testing
  - Performance test suites, regular large scale tests