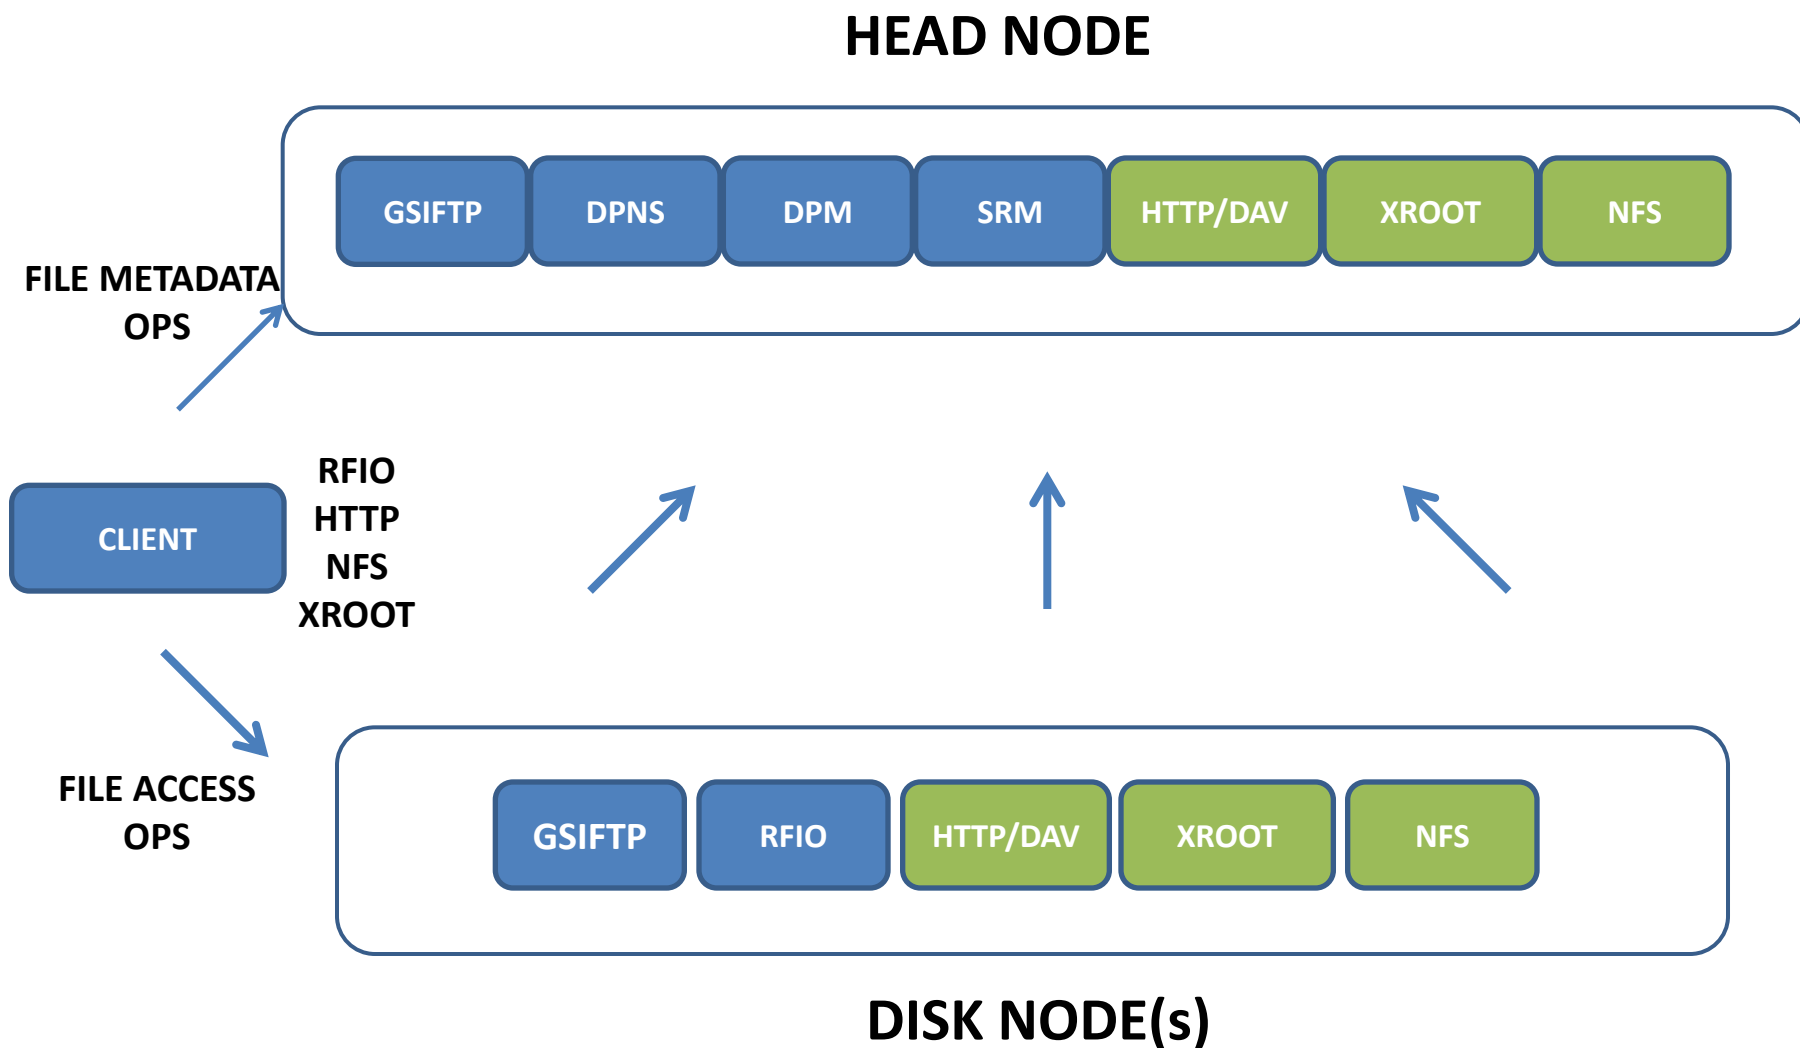



DPM interfaces

Oliver Keeble, CERN
On behalf of the DPM team

DPM daemons



- The following are not receiving further development
 - RFIO
- The following will get or have got a bit
 - SRM (to include support for new backends)
 - Gridftp (redirects, see later)
 - DPM native (now has sync get)
 - DPM NS (thread pool)
- The action is here
 - HTTP
 - Xroot
 - NFSv4.1






- We are pushing HTTP/DAV as the single protocol solution
 - It works for WAN transfers
 - 3rd party copy, multistream solutions exist for DPM & dCache
 - It works for LAN random access
 - See results from Martin Hellmich
 - It's OK for WAN random access
 - We're tuning
 - Federations have been demonstrated
 - Allows a GET of an LFN or "SURL" to transparently redirect to the disk server
 - Avoiding SRM



HTTP/DAV



Success rate: 100% (4/4)
Start time: Tue Oct 02 15:31:41 2012
End time: Tue Oct 02 15:31:57 2012
Duration: 0h 0m 15s

dav-basic [SUCCESS]

dav-integrity [SUCCESS]

dav-rmdir [SUCCESS]

dav-spacetest [SUCCESS]

```
# /usr/share/lcgdm/tests/functional/dav/dav-spacetest
```

Exit code: 0

```
Space reserved with 209715200 bytes
100M file created
[201:UPLOAD] /tmp/dav-reserved.tmp to /dpm/cern.ch/home/testers.eu-emi.eu/reserved.1349184712?spacetest=dav-reserved
Current free space: 104857600 < 209715200
[507:UPLOAD] /tmp/dav-reserved.tmp to /dpm/cern.ch/home/testers.eu-emi.eu/reserved.1349184712.2?spacetest=dav-reserved
Current free space: 104857600 < 209715200
Second attempt failed with 'Space unavailable' (507) as expected
[204:DELETE] /dpm/cern.ch/home/testers.eu-emi.eu/reserved.1349184712
[404:DELETE] /dpm/cern.ch/home/testers.eu-emi.eu/reserved.1349184712.2
Space released
[400:UPLOAD] /tmp/dav-reserved.tmp to /dpm/cern.ch/home/testers.eu-emi.eu/reserved.1349184712?spacetest=dav-reserved
Third attempt failed with 'Bad request' (400) as expected
-TEST PASSED-
```


nagios [SUCCESS]

Success rate: 100% (1/1)



GOAL: functional and performance improvements.

- Performance is of primary importance
 - Profit from changes in DPM
- Allow xrootd clients to access and write files to DPM
 - Use X.509 authentication, with VOMS extensions
 - DPM does the usual authorization
 - Optionally allow the ALICE “tkauthz” mechanism
- Mixed multi-VO setups
 - A single DPM can serve more than one VO
- Support for per-VO xrootd federations
 - Following the evolution of the ATLAS LFC-based name translations
- Tested in the FAX project in Atlas
 - Successful integration of LFC via name2name plugin
 - Will have production release in DPM 1.8.4



- Same for PUT
- This is not ideal!



- GFD20 (OGF standard)

3.2. Commands

3.2.1. Striped Passive (SPAS)

This extension is used to establish a vector of data socket listeners for each stripe of the data. To simplify interaction with the parallel data transfer extensions, the SPAS MUST only be done on a

allcock@mcs.anl.gov

7

GWD-R (Recommendation)

April 2003

control connection when the data is to be stored onto the file space served by that control connection. The SPAS command requests the FTP server to "listen" on a data port (which is not the default data port) and to wait for one or more data connections, rather than initiating a connection upon receipt of a transfer command. The response to this command includes a list of host and port addresses the server is listening on. This command MUST always be used in conjunction with the extended block mode.

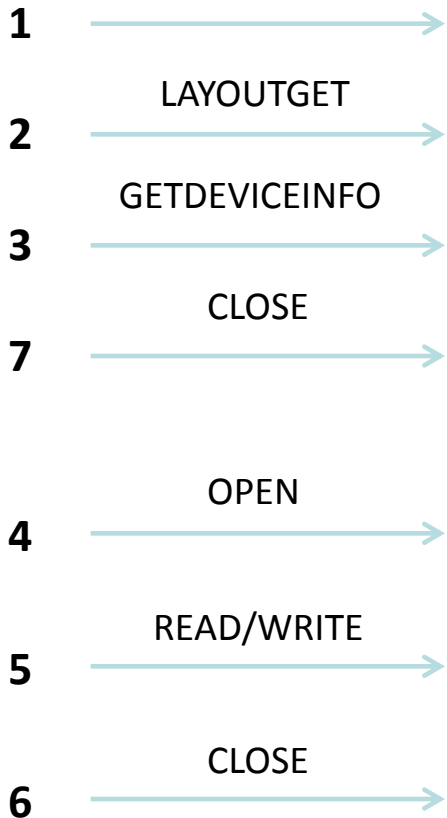


- Plan would be to use “Striped Passive” with a single stripe, to profit purely from the redirection
- Ricardo checked current implementations
 - Note, gridftp implementations are typically globus gridftp + DSI plugin for each backend
 - This has to be supported in the plugin
- He found
 - Only the HPSS plugin appears to support this
 - Unable to verify if it works
 - Implementation is non trivial
 - Spacetokens?
- We can investigate this further for DPM



NFSv4.1

CLIENT



METADATA SERVER

DATA SERVER



- It's a standard way of doing what we already do
 - Direct access to data nodes
 - Redundancy, parallelism
 - WAN friendly features
 - Kernel caching advantages
- Available in latest DPM releases r/o
 - Based on ganesha server
 - r/w by end of EMI (Apr 2013)
- Under continual testing at ASGC Taipei
- Only kerberos auth
 - Still looking for the best way to do X509



- As a disk-only solution DPM is a good place to explore SRM alternatives
- New HTTP, xroot and NFSv4.1 interfaces offer much improved access
 - And alternatives to some aspects of SRM:
 - Getturl
 - Write to spacetoken
 - Checksums