



DPM Workshop Summary

Oliver Keeble
On behalf of team DPM
IT/SDC

12/11/2014



5th DPM workshop

- 9-10 October, Naples
- <https://indico.cern.ch/event/324705/>
- Around 20 participants and 7 remote attendees
- All major DPM infrastructures represented:
 - Australia, France, Italy, Poland, Taiwan, UK
- Project news, site reports, prioritisation for 2015
- THANKS
 - INFN, Alessandra Doria and Alessandro De Salvo

DPM Collaboration

- Quick review of the DPM Collaboration activities
 - A number of projects have pledged effort to DPM
 - Most testing now undertaken by partners
 - Early adoption of new features
 - Participation in staged rollout
 - Provision of testbeds for nightly tests
 - Admin tools, rebalancing utility (GridPP)

Site Reports

- A very brief flavour of the different reports:
- ASGC
 - Lots of work to make SRM perform
- UK
 - Interest in different backends (Lustre, HDFS)
 - Argus integration
 - Need ways to control load spikes
 - Interested in non-SRM operation

Site Reports

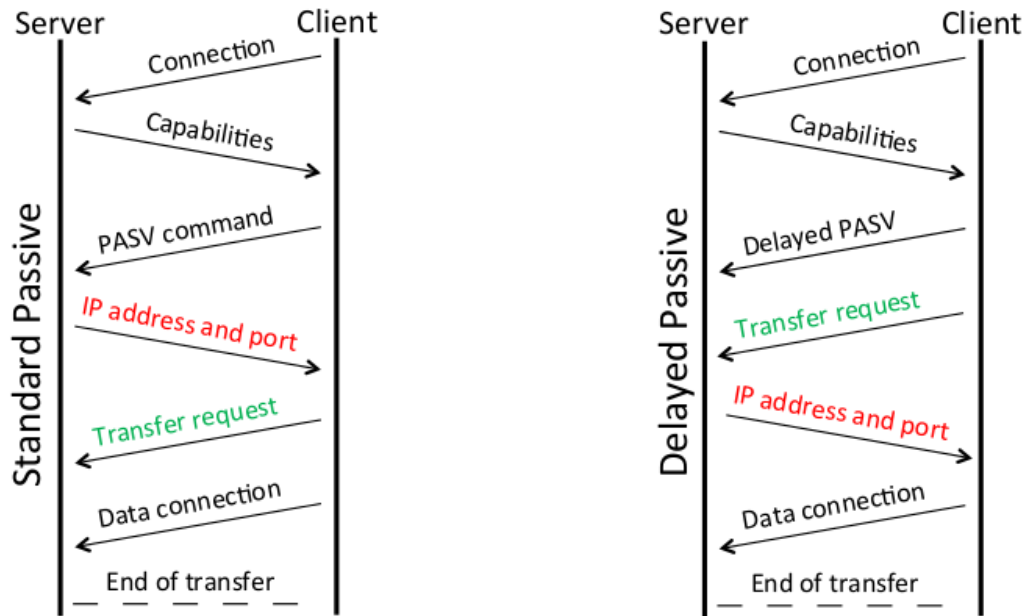
- Italy
 - Investigating novel deployment scenarios
 - High Availability
 - Distributed T2 prototype (Napoli-Roma)
 - Federated access (inc StoRM)
- France
 - Full AAA/FAX participation
 - Struggles with draining and balancing
 - Quattor based, looking at hybrid with puppet
- Australia
 - Part of Canadian HTTP “dynafed” federation
 - Also run DPM for Belle II

1.8.9

- DPM 1.8.9 has just been released
 - i/o monitoring “a la xrootd”
 - New logging capabilities
 - Performance improvements
 - End of YAİM support, full puppet support
 - Infosys improvements
 - MySQL/memcache improvements
 - ARGUS support re-enabled

GridFTP redirection

What exactly is Delayed in Passive mode?




GridFTP updates

- Latest releases support a redirecting mode of operation
 - Previously, data was tunnelled within the system
 - Large potential performance benefits for gridftp-only transfers
 - Requires activation via configuration
 - And monitoring of client activity
 - Old clients can be inefficient – please move to gfal2!!
- Presentation focused on motivating some “early adopters” to validate the system

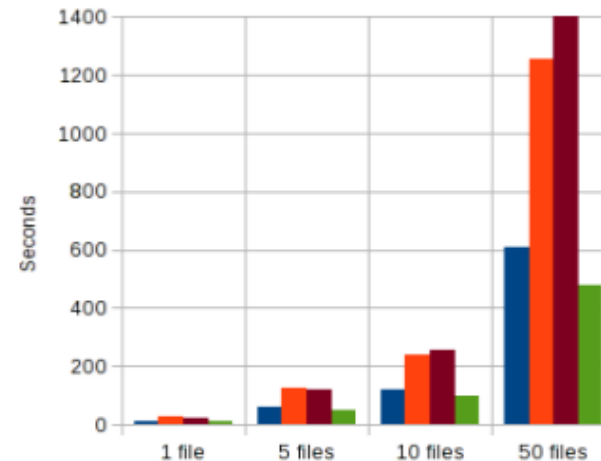
Performance study

Sequential transfer of files

Small file (40MB) transfer performance



Large file (1GB) transfer performance



- SRM + GridFTP (lcg-cp -b)
- GridFTP + int. RFIO (lcg-cp -b)
- GridFTP + int. RFIO (globus-url-copy -dp)
- GridFTP redir. (globus-url-copy -dp)



IT-SDC

GridFTP 2 – the road to redirection

9/10/2014

13

/tmp/kiryanov-gridftp.pdf

[13/16]



IT-SDC

dpm-xrootd federation

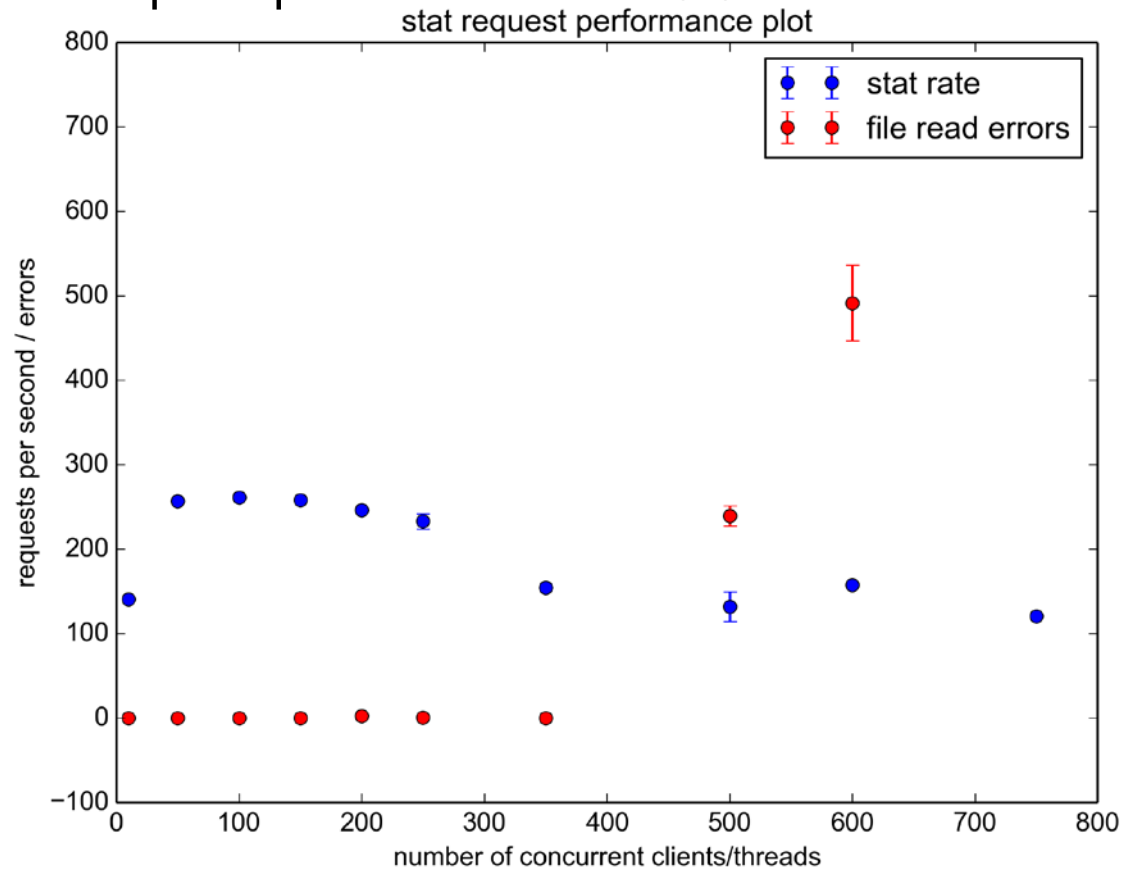
- Can be configured with yaim (or manually) or with puppet. See:
 - <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Xroot/Setup>
 - <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Xroot/ManualSetup>
 - <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Xroot/PuppetSetup>
- **vomsxrd** should be installed
 - To give full voms information for authorisation
 - VO name in exported monitoring data

CMS File opening and reading tests

- Tests emulate CMS jobs running at CMS sites, choosing the site through regional redirectors
- CMS target for initial tests:
 - File-opening test: access total rate of 100 Hz at a site
 - File-reading test: 600 jobs reading average rate of 2,5 MB every 10 s at a site → reading total rate of 150MB/s
- these numbers come from internal CMS analysis, based on historical figures
- tests reveal sites that need further optimization and possible improvements
- these are not meant to be a stress test for the site


Nightly perf tests

- Example: HTTP peak performance in 1.8.8

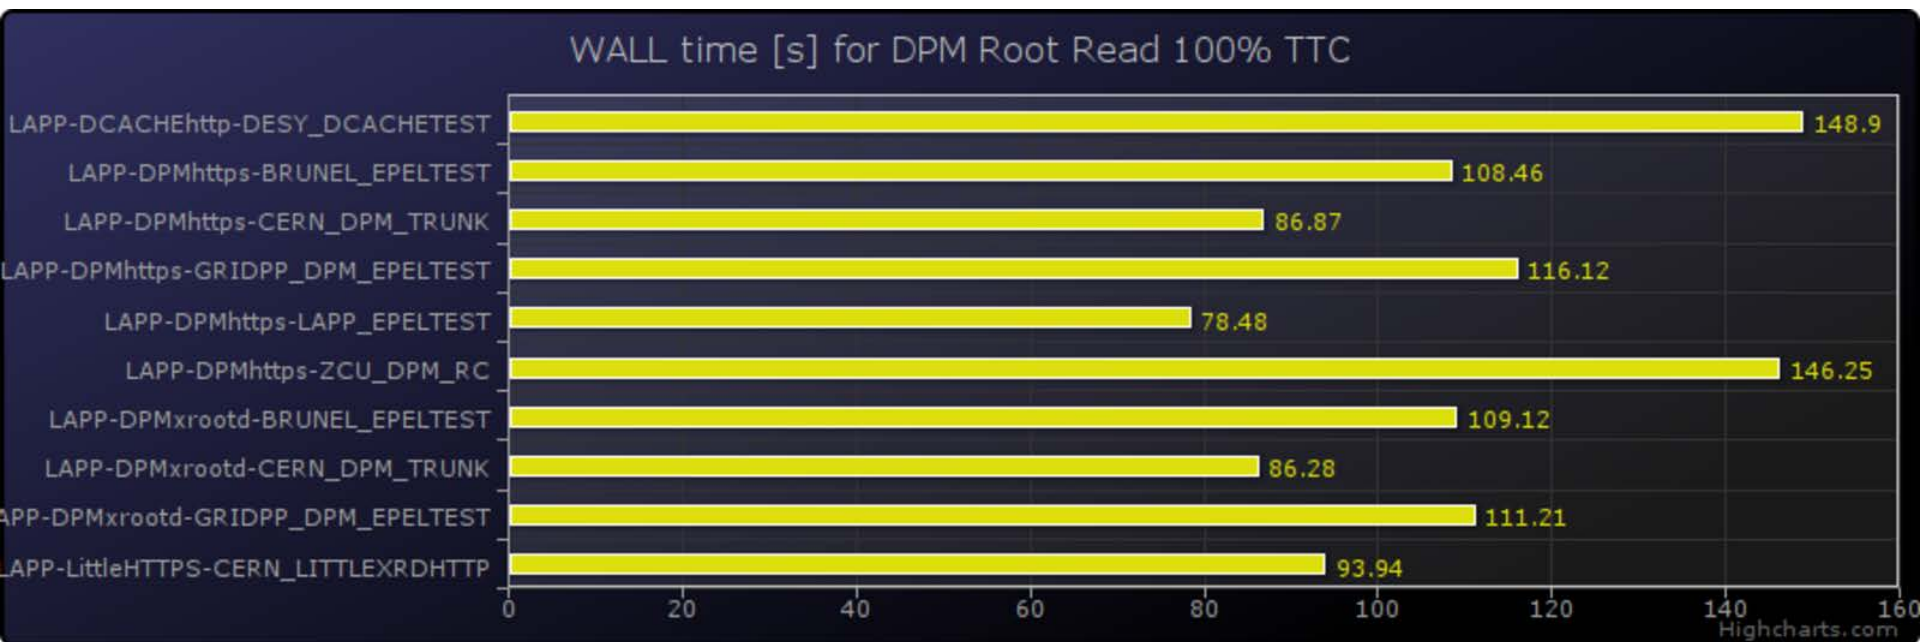


Nightly perf tests

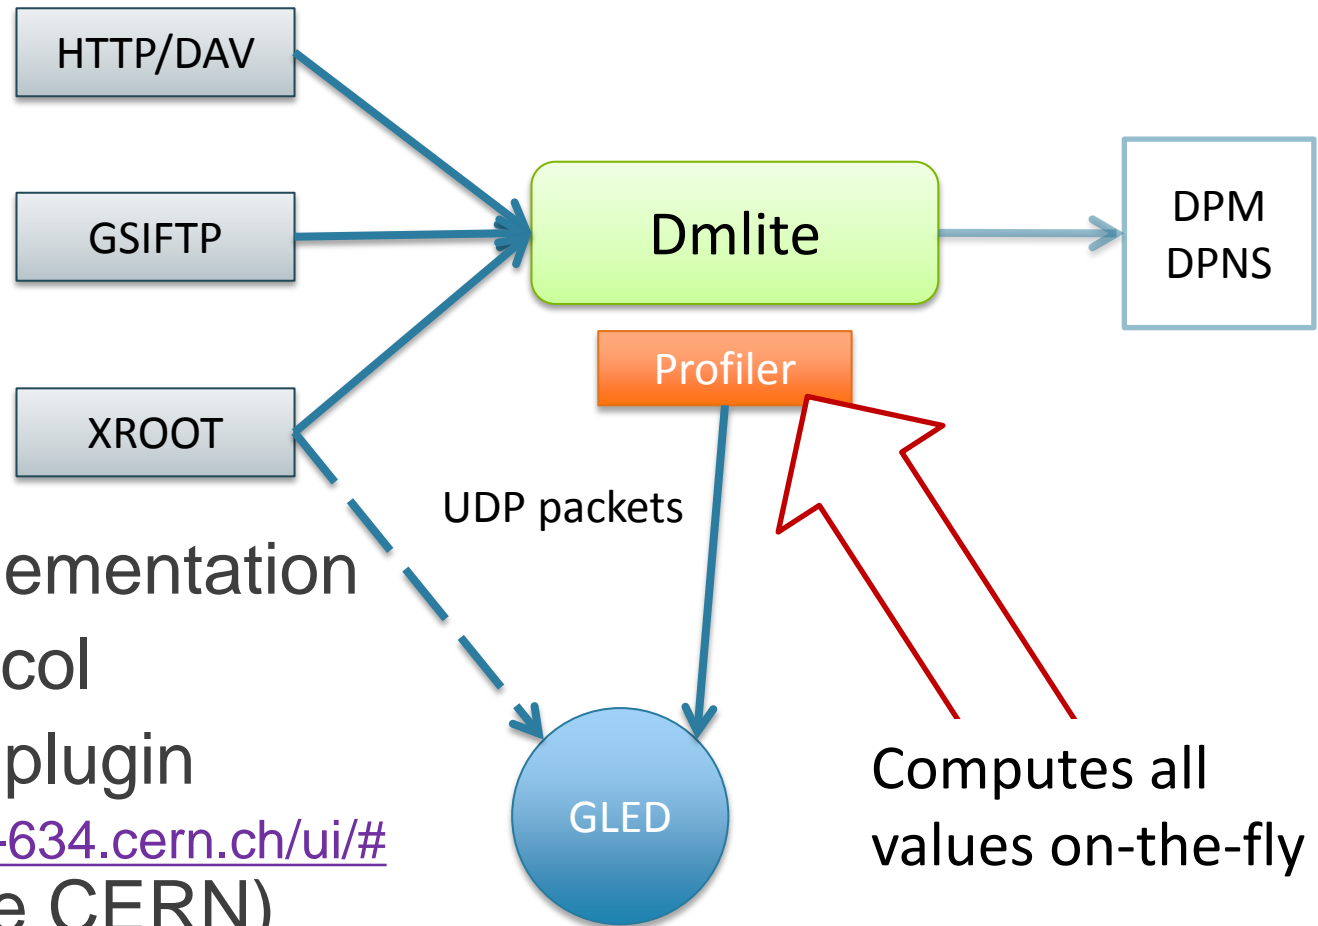
- Example: HTTP peak performance in 1.8.9 ... a 20X increase



Continuous testing with Hammercloud



Monitoring in dmlite



- Single implementation
- Multi-Protocol
- Separated plugin
- <http://dashb-ai-634.cern.ch/ui/#>
(from inside CERN)

HTTP Federation

starting PoC with DPM sites

Group by dst/clients that accessed *cern.ch*



DPM Configuration evolution (as presented during the previous Workshop)



YAIM



IT-SDC

Using Puppet for DPM

10/10/2014

3

/tmp/Using Puppet for DPM.pdf

[3/40]



IT-SDC

DPM Workshop Summary, GDB Nov 2014

17

Configuration

- YAIM support is now deprecated.
- Puppet is fully supported, all new configuration logic will appear there
- Usable without a “puppet master”
- All necessary config available through puppetforge
 - Also uses some CERN modules
- Full example manifests documented

Future

- Minimise maintenance costs
 - Remove/minimise legacy code path
 - -> “dmliteSE”
- SRM-less operation
 - Space reporting
- File placement, rebalancing, drain
- Incorporation of alternative backends

Space reporting : Proposal

- Agreement in principle with **dCache** and **StoRM**.
- RFC 4331 “Quota and Size Properties for Distributed Authoring and Versioning (DAV) Collections”
- Adds two new DAV props:
 - DAV:quota-available-bytes
 - DAV:quota-used-bytes