

DPM WORKSHOP

Highlights

Frederique Chollet
Andrea Sartirana

❖ Naples, 9-10 Oct 2014:

- <http://indico.cern.ch/event/324705/> ;
- **thanks** to the **INFN** for the hospitality;

❖ ~20 participants

- **Australia, France, Italy, Poland, Taiwan, UK;**
- France represented by F. Chollet, A. Sartirana (+ M. Jouvin remotely);

❖ just a **brief summary**. See also:

- <https://indico.cern.ch/event/272779/> (Keeble).

the collaboration

❖ DPM now run as a **community effort**:

- **CERN** takes care of the **core development**;
- **partners**: **complementary skills/resources**

It was not possible to put this proposition (web admin interface) into practice.

GRIF_{IRFU|LLR}
active on puppet.
QWG care of quattor.

GRIF_LAL, LAPP
+ GRIF_LLRL contrib.
to WLCG middleware
validation.

Area	Partner
DMLITE/VFS	CESNET
Admin tools, rebalancing	GridPP
I/O Monitoring	Japan
Web admin interface	France
Support (beyond local responsibilities, not necessarily on rota)	GridPP
Documentation & web presence	GridPP, ALL
Configuration, puppet, YAIM transition	France, GridPP
Permanent pilot testing (EPEL-test) – for validation of all updates	GridPP, France, Italy
Ad-hoc testing: HTTP/DAV	GridPP
Ad-hoc testing: S3 validation	GridPP
Ad-hoc testing: HTTP federations	Italy, GridPP
Ad-hoc testing: IPv6	GridPP
Ad-hoc testing: GridFTP redirection	Needed
Hammercloud & I/O optimisation	GridPP, ASGC
All activities not mentioned above, see note	CERN

echos from sites



- Struggling with **SRM perfs.**



- Fully **AAA/FAX** enabled;
- **drain/load bal.** pains;
- mostly **Quattor**. Some **puppet** and **hybrid** sites.



- Open Stack **High Avail.**;
- **distr. T2 proto** (NA-RO)
- **http/WebDav fed.** (inc **StoRM**) .

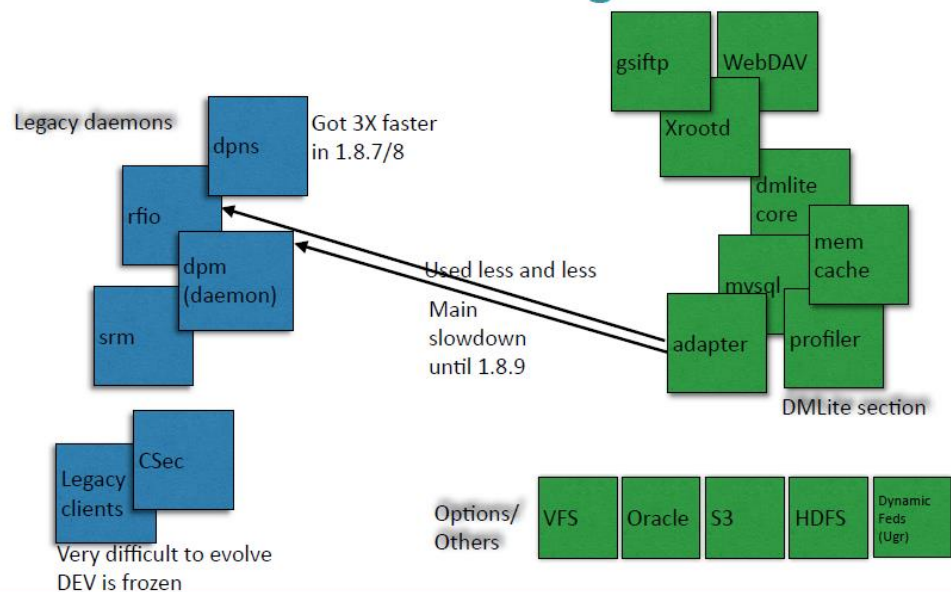


- Interest in **diff. backend**:
Lustre, GPFS, HDFS, Cloud;
- **argus** integration;
- need **load spikes mgmt**;
- interest on **non-SRM ops.**



- Part of the **Canada HTTP dynafed federation**;
- also run **DPM on Belle II**

the dev status



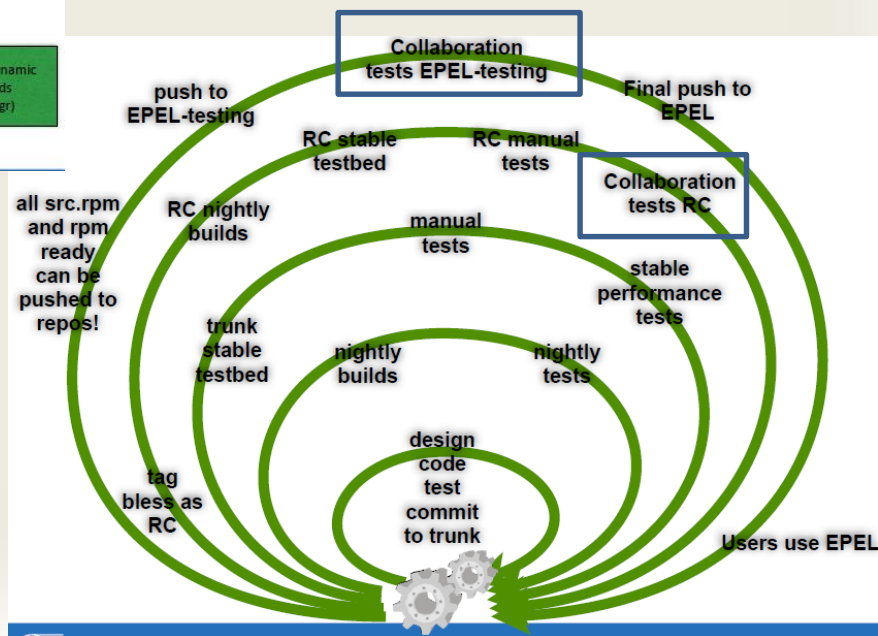
... DPM is a set of **components blessed as well working together** plus a standard **configuration** given by **puppet** templates ...

F. Furano



... Our **continuous build and testing** workflow proved to be very **reliable** and helped us a lot to use well our time...

F. Furano



❖ Released on **05/11** in **epel**:

- lcgdm 1.8.9, dmlite 0.7.1, dpm-dsi 1.9.3, dpm-xroot 3.3.6, webdav 0.15 ;
- @ FR sites: GRIF_{IRFU|LAL|LLR},... ;

❖ **not yet** in **WLCGBaselineVersions**

- delayed until new SAM probe available **(15/12)**;

❖ highlights:

- **performances**: up to a x10 wrt 1.8.8;
- new features: **logging** and **xrootd-like mon**;
- **official puppet** modules;
- other: dpm-listspaces fix/argus poller emi3.

performances in 1.8.9

❖ **A lot of work** on performances:

- **memcached** is now part of the standard install ;
- **removed** most of **interactions** with **dpm/dpns**;
- up to **x10 peak perf** wrt 1.8.8 in opening files;

❖ triggered by **CMS-AAA tests**:

- round of test ran by **F. Fanzago**;
- target: **100Hz** file open / **150MB/s** file read;
- **precious feedback** from sites.

new in 1.8.9: logging

❖ New **DMLite logging** subsystem:

- log **by level** log **by component**;
- final filter through syslog;
- many **log lines** added to **all components**;
- in the config file:

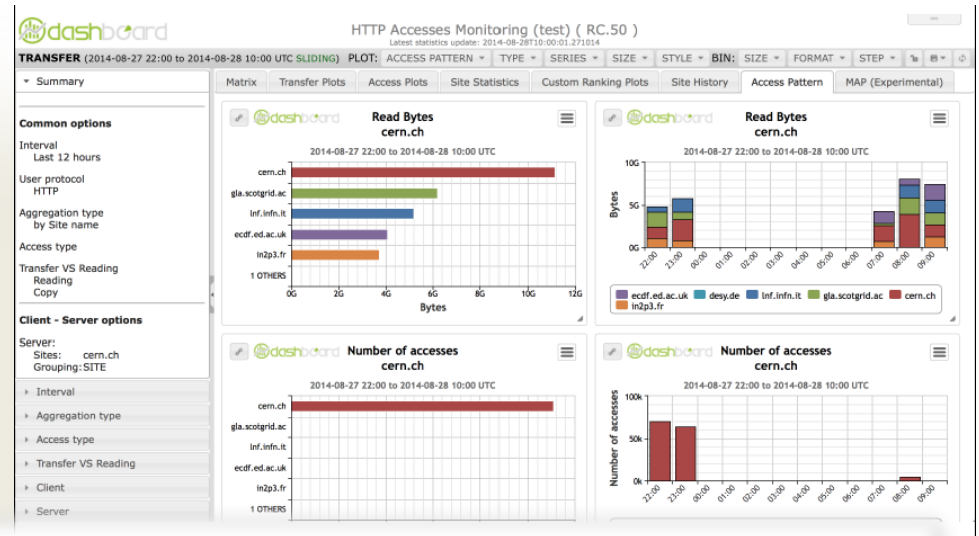
```
■LogLevel=[0..5]
```

```
■Log <componentname>
```

```
2014-09-19T14:22:14.740239+02:00 lxfsra10a01 httpd.event[12384]: dmlite Memcache getValFromMemcachedKey : Exiting with value from memca
2014-09-19T14:22:14.740249+02:00 lxfsra10a01 httpd.event[12384]: dmlite Profiler extendedStat : path: /dpm/cern.ch/home, follow: 1
2014-09-19T14:22:14.740265+02:00 lxfsra10a01 httpd.event[12384]: dmlite Mysql extendedStat : parent:0 name:/
2014-09-19T14:22:14.740468+02:00 lxfsra10a01 httpd.event[12384]: dmlite Mysql execute : Executing: 0x7f75ec046f40 nParams_: 2
2014-09-19T14:22:14.740550+02:00 lxfsra10a01 httpd.event[12384]: dmlite Mysql execute : Executed: 0x7f75ec046f40 nParams_: 2 nRows:1844674
2014-09-19T14:22:14.740573+02:00 lxfsra10a01 httpd.event[12384]: dmlite Mysql extendedStat : Exiting parent:0 name:/sql
```


- Add to your manifest:

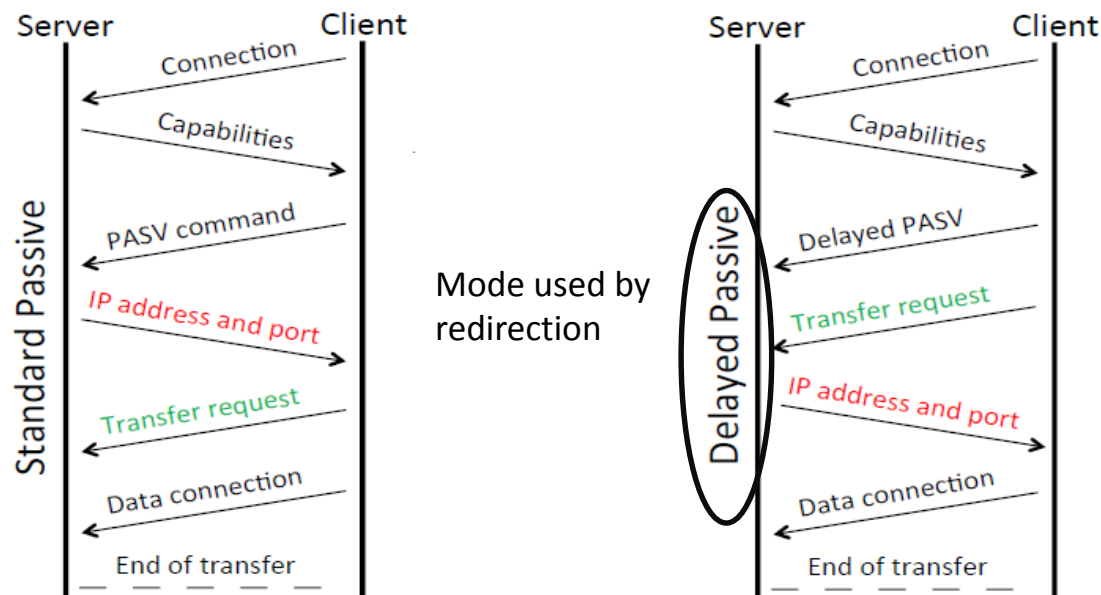
```
class{"dmlite::plugins::profiler":
  collectors => ["rndgled.cern.ch:9930"],
  auth      => "userdn"
}
```



- ❖ Dmlite-plugin profiler plugin:
 - wraps **dmlite** api calls;
 - sends info to a **GLED**;
 - implements **real-time multiVO multiprotocol xrootd-like IO monitoring**;
- ❖ **integration of HTTP to monitoring dashboard in progress.**

gridftp redirection

- Allows for **SRM-less** gftp transfers;
- <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/GridFTP>
- **not all clients** support the DP mode needed
 - **old clients** will trigger **rfio** background **copies**;
 - **FTS3** and **gfal-2** are ok.



Available since
1.8.8

Currently
testing: LAPP,
GRID_LLRL and
Edinburgh.

looking forward

❖ **Replace RFIO** with WebDav:

- **Davix** seems to have all needed functionalities;

❖ **review** spacetoken and **pools**:

- **remove pools** as we know them;
- **space management on top level directories**;
- **disks labelling**;

❖ space reporting : **srm-less** operation

- cover srm functionality within the namespace ;
- proposal and **consensus with dCache and StoRM** on metadata report (webdav RFC) .

since the workshop

❖ **Gridftp logging** too verbose

- `fix provided;`

❖ **reconsid. of metapackage** for handling deps

- metapackage **added to epel**, trying to avoid "synchronization pb" with EMI metapackage;

❖ **rebalancing / draining**

- in the forum, **sites stuck with rfio** (little improvement in 1.8.9 fix which prevents a hang but still bug in rfiod);
- preference to **redo the tool** using the recently API **using webdav and gridftp.**

❖ DPM community is **very active**:

- sites and devs seem to be making a great job;

❖ **1.8.9** released:

- important **performance upgrade**;
- some **new stuff**: logging and mon;

❖ other **new** stuff:

- **gridftp redirection**;
- ...and **more to come**: rfio/pools/srm bye bye;

❖ **French** contribution:

- **testbeds** running at LAPP, GRIF_{LAL|LLR};
- active on **puppet/quattor** conf - GRIF_* .