

DPM and the upgrade TF

March 2019

Fabrizio Furano, Oliver Keeble, Andrea Manzi

DPM – General direction

- DPM codebase - pretty stable, no revolutions in the code since 2 years
 - “New” (since 2017) codebase (DOME): XrdHTTP implementing a REST interface
 - Old codebase (lcg-dm) is not central anymore, it’s optional and serves only SRM (SRM stays functional, becomes optional)
- DPMs can be configured in the “legacy” or in the “dome” flavour
 - Both flavours support SRM
- Main topics are refinements, support and upgrades.
- We will stop supporting the old codebase in May 2019 (announced many times since one year ago)

LCGDM support from 01/Jun/2019

- From **1st of June, 2019** our standard LCGDM support answer will be **“there is an alternative: upgrade to DOME flavour, please”**
 - That affects: dpns, dpmdaemon, rfio, CSec, dmlite::Adapter, SRM
- LCGDM will stay in EPEL as long as it compiles untouched in Rawhide (EPEL rules will remove it the day it breaks)
- It's pure C, hence that can be even years, we don't give limits

Main points distilled

- SE sizes and load increase. The DOME flavour comfortably copes, the older components ... less
- WLCG Space reporting (SRR): available only with the updated system (as it's based on directories, not spacetokens)
- Full checksum support (also for xrootd, with TPC) needs the updated system in DOME mode

WLCG Storage Reporting Record (SRR)

- DPM \geq 1.10 can produce the WLCG SRR
 - Since it pioneered this, DPM has been chosen for a “pilot” WLCG Task Force in Q4/2018
- The “DOME mode” (= modern codebase) must be enabled
- The TF is progressing, the idea is that after the pilot phase all the DPM WLCG sites should be proposed to upgrade
- Challenge: many obsolete gridftp clients (lcg-utils) are still in the wild, and they are suboptimal with gridftp2 (gridftp redirection)
- Challenge: often the gridftp2 redirection support is confused with enabling DOME. This is a communication challenge.

DPM upgrade TF - Pilot phase

- In the context of WLCG some volunteers joined the pilot phase for the DPM upgrade TF
- We wish to thank them very much for their contributions, help and advice.
 - <https://twiki.cern.ch/twiki/bin/view/LCG/DPMupgrade>
- Most of them upgraded/reconfigured just fine
- As expected, the activity triggered discussions, improvements and fixes to the docs and to the components
- What follows is a short summary

DPM upgrade TF - Pilot phase

- Petr Vokac contributed polishing one the migration tools, quite a few fixes and productive discussions/reports
- Cross-protocol activities triggered improvements
 - When writing through SRM and deleting with WebDAV, the legacy DPM daemon was not adjusting its internal free space counters [[LCGDM-2752](#)]
 - Consequences: mismatch in the dpm-listspaces vs dome numbers, plus the DPM daemon might erroneously think that the disk was full
 - This was fixed in the DPM daemon (yes, the good old one). v1.12 keeps the values aligned, no matter what.
 - A.Doria (INFN-NA) had reported this

DPM upgrade TF - Pilot phase

- DPM daemon was allowing deletions in r/o filesystems, Dome instead was denying. Changed the behaviour in Dome. [[LCGDM-2735](#)]
- The very high load of Brunel triggered two issues
 - A bug affecting stability (fixed) [[LCGDM-2753](#)]
 - An incorrect default value for some parameters (sizes of internal DOME resource pools/ queues. This was piling up checksum requests). Matter of configuration, yet caused clogging. Many thanks to Raul Lopes for his help understanding the issue. Also Prague was seeing this. [[LCGDM-2745](#)]
- Brunel also showed that the GridFTP redirection is more demanding on the core components (dpns/dpm/rfio). We improved the docs to advise to enable it only after Dome is enabled
- Another investigation is pending in Brunel about 0-sized files and the behaviour with gridftp uploads that abort for external reasons. Still ongoing.
- The last fix was a bug affecting performance slightly: Database connections not reused [[LCGDM-2754](#)]

After the pilot phase

- The pilot phase lasted longer than the optimistic provisions, yet it's converging. The instructions have been polished:
 - <https://twiki.cern.ch/twiki/bin/view/DPM/DpmSetup>
- We were waiting for 4-5 sites, we have 6-8
 - <https://twiki.cern.ch/twiki/bin/view/LCG/DPMupgrade>
- We need to release v1.12, with the fixes. The EPEL release takes 2-3 weeks
- We think that the future steps should be discussed within WLCG. A few options:
 - A second wave of guided volunteers
 - Do nothing and let the sites decide when they need to support SRR, full checksumming, higher load and if they want to run unsupported software in production

DPM Workshop (Bern 13-14 June 2019)

- We will discuss these things also at the DPM workshop (agenda coming soon)
- <https://indico.cern.ch/event/776832/>