

DPM and Dynafed 2020

Fabrizio Furano
Nov2020

Intro - What's in this presentation

- Describe the path that led to DPM and Dynafed
- Understanding what constitutes the effort of maintaining these systems
- Hint at the milestones that the systems will have to face

- Promote once more the “open source” character of the modern DPM/Dynafed components
 - ... which is the best chance for longer life and support with a healthy coordinated community

Modern DPM and Dynafed: DMLite

- in 2010-2012 the CERN “Software for Distributed Computing” group started developing the DMLite framework and libs. The activities were part of the EU EMI project
- DMLite is an abstract library to build systems for managing storage
- DMLite spurred in 2010-2012 the writing of Dynafed
- The HTTP client Davix (used also by ROOT) came from Dynafed, and found an important place in the present and future of computing at LHC scale
- DMLite was a shy attempt to incrementally rewrite DPM, which was facing code obsolescence at the horizon
- This focused in 2017-2019 with the DOME daemon
 - Made possible to implement new features, fix historical troubles and obsolescence and improve dramatically robustness and scalability of DPM

DMLite → DPM and Dynafed

- DPM: traditional “GRID storage element” system based on a DB catalog and the concept of “disk pools” and a full-featured redirector node (“head node”)
- Dynafed: minimalistic, very flexible HTTP(s) dynamic redirector
 - We can see it as a DPM without DOME and not needing a static catalog
 - The catalog is built on-the-fly in memory
 - We used to call this “storage federations”, yet it became more than this when someone realised it works very well with S3 and caches in general
- Both systems support things like TPC, X509, VOMS, Macaroons, OpenID-Connect
- DPM comes with a puppet-based pre-made do-it-all config

DPM Status - November 2020

- Total disk space: 96PB
- Number of instances: 88 (BDII says 71 in 59 sites)
- The old components (dpnsdaemon, srm, dpm-daemon, rfio...) are deprecated since September 2019
 - No urgency, they will stay in the EPEL7 repos, simply not being ported to EPEL8 and not updated with newer version numbers
- Well up to date with the features needed in WLCG (e.g. Macaroons and OpenID-Connect, cross-protocol checksums, TPC, easier multi-site, pools as caches)
- The DPM upgrade TF has tracked/promoted the upgrade progress and the enabling of the new components (DOME) and of the WLCG Storage Reporting Record
 - <https://twiki.cern.ch/twiki/bin/viewauth/LCG/DPMupgrade>
- Roadmap of stability: the tech goals have been reached, hence no revolutions are foreseen in the DPM development
- Given the stability of the platform, the dev deltas that we see will mostly be little fixes and polishments, e.g. CLI commands, security fixes, etc.

Dynafed status - November 2020

- Relevant deployments:
 - BOINC CMS@home to store files in S3 at CERN
 - Canadian S3-based cloud (with redirectors at CERN)
 - ECHO at RAL
- Differently from DPM, we don't track Dynafed installations
- Can be configured with TPC and OpenID-Connect
- Particularly flexible, well documented authorization subsystem

Maintenance main ingredients

- Facilities at CERN:
 - GitLab repository group: <https://gitlab.cern.ch/lcgdm>
 - Jenkins build system: <https://jenkins-lcgdm.web.cern.ch/>
 - Semi-automatic test procedures
 - Small testbed
 - Seriously testing the development releases of a plethora of other components as a side effect, e.g. xrootd, voms, davix, gfal, etc.
- Precise, simple, written down test/release workflow
 - Pretty quick turnaround in case of urgency
 - A few brave sysadmins also help checking things in their sites
 - EPEL is the last step of the release
- In the last years the development effort has been negligible with respect to the running these necessary steps

Publish to Fedora/EPEL8

- The code does not need fixes, or maybe just minimal things
- Petr Vokac has already contributed some needed bits
- The specfiles might need juggling
- The transition to Python3 is not fully complete yet
- Full status: <https://twiki.cern.ch/twiki/bin/view/DPM/DpmEpel>
- The effort would be also about upgrading the build/test system and the testbeds dpmhead-trunk and dpmhead-rc
 - Xlation: “fight with specfile, Openstack, Jenkins and Puppet”
 - Hey! There can be no releases without builds/tests!
- CERN will not contribute directly the EPEL8 porting, as per statement
 - <https://indico.cern.ch/event/813745/contributions/3766117/note/>
 - This means that somebody could be accompanied to do it



Summary: Roles involved

- Which roles are necessary for a project in maintenance mode?
 1. Community (utilises, discusses and contributes)
 2. Partial time dev (reviews contribs and contributes himself sometimes)
 3. QA manager (maintains the test clusters dpmhead-*, e.g. for DOMA tests or nightly internal ones, reviews the tests results)
 4. Build/test system maintenance (Jenkins+OpenStack+Puppet take some time to maintain)
 5. Epel master (for the final release steps)
- For DPM, points 2, 3, 4 have been inherited by Fabrizio
 - Point 5 will soon be vacant

AAAAAA - EPEL packager wanted

- Oliver Keeble moved to other things, and can't anymore guarantee this role
- Fabrizio never was an EPEL packager
- In absence of a packager, we can't push to EPEL even small fixes (e.g. security) from now on
 - And coincidence wanted that we had a security fix in the pipeline now (thanks Oliver and Petr for making possible v1.14.2)
- The srpms to push are just two
 - dmlite (contains all DPM and dmlite)
 - dynafed (depends on dmlite)
- These packages will be orphaned in the next few weeks (by Fedora policy recommendations)
- (And maybe this is the right time to update these things in UMD, to make sure that a coherent repo stays)

Status

- Very stable service and codebase
- The code is in Gitlab, totally open
- For DPM, the scalability range is pacifically over the “tens of PBs”, the performance is more than adequate for this. Dynafed is one level up, depending on the usage, being a more lightweight-featured system
- Thanks for the contribs :-) :-)
- Will stay in EPEL7 until EPEL7 exists
- We see only reasons for DPM/Dynafed to work well in the next few years, also depending to the vitality and interest of the community

- Oliver Keeble has been the previous EPEL packager
- No plans so far to push to EPEL8, a non-CERN packager could do it

Conclusions

- Modern DPMs profit from a healthy technical platform. This includes Dynafed as a successful spin-off
- The system accommodates all the current requirements (including TPC and bearer tokens) and is technically well-placed to accommodate future ones that may come
- Discussions are ongoing about the directions for the DPM support, both from the point of view of the sites, EGI and of the WLCG Operations team
- Overall, the DPM project is more than 15 years old and was incrementally refurbished and put up to date with the modern requirements
- We all put hope into the cooperation of the DPM/Dynafed community, WLCG, CERN and EGI so that such an important asset benefits sites for many more years
- Now... it's time for free questions and clinic