



DPM Migration campaign status

Renato Santana / Alessandro Paolini

EGI Foundation

16-02-2023

OMB meeting

Introduction

- In [Sept 2022 tickets](#) were opened to sites providing DPM endpoints
- Sites have been encouraged to start the migration to a different storage element since the process will take time
 - Choosing the new storage solution depends on the expertise/experience of the sites and on the needs of the supported Vos
 - Currently available in UMD: dCache, Storm, EOS, XrootD.
- DPM team provided a script and a guide for the migration to dCache:
 - <https://twiki.cern.ch/twiki/bin/view/DPM/DpmDCache>
- Collecting migration plans
- Following-up the tickets along with DPM team to provide support in case of issues
- Migration to dCache should be completed by **June 2023**.

DPM Decommission and migration campaign

Status (link to [summary overview](#))

Number of tickets	
In progress	27
On hold	16
Not replied	3
Solved	9
Total	55

Planned completion date	
By 2022	1
Beginning of 2023	5
By Feb 2023	5
By Q1 2023	8
By Q2 2023	12
undefined	11

Chosen technology	
Migration to dCache	28
Migration to EOS	6
Migration to xroot/ceph	3
Migration to XrootD	1
Migration to Xcache	1
Migration to Dynafed	2
Not yet decided/no clear plan	11
Decommissioning SE or site	3

Issues encountered (I)

Found stuck replica – [GGUS 158796](#) (GRIF – NGI_FR)

1. “found stuck replica” means that file transfer did not finish correctly (in the `dmlite-shell -e 'info /path/to/file'` you would see status `kBeingPopulated`) – e.g. there was a race-condition in the DPM `gsiftp` implementation in the past that could lead to this state or also if disk-node suddenly died you would see also these kind of inconsistent state.
2. Anyway, it is safe to remove “found stuck replica” files, because their transfer to the DPM did not finish correctly and they were not accessible for reading via DPM protocols. (In the migration documentation there are commands to clean all “dirty/inconsistent transfers”).

Issues encountered (II)

[GGUS 158832](#) (UKI-SCOTGRID-ECDF)

1. The migration tools had to be adjusted to work in an environment where we have to use custom ssh ports. (I have a fix for reading the ssh configs which I need to work out where to contribute)
 - nonstandard ssh ports – that script should at least log some reasonable exception (the used SSH config is rare)
2. The initial migration failed due to one of the quota tokens in DPM containing more data than was in one of the spaces which caused problems at the db import. (This was fixed by starting again on the import after modifying the space allocation by hand in a .csv file)
 - added warning in case spacetoken size is smaller than sum of file size from given spacetoken (but next dmlite release uncertain)
3. ATLAS is now picking up files not found. Unfortunately a) they didn't exist in the site DPM, b) were not identified on disk as damaged or dark-data detected in the migration from DPM and c) appear not to have been written since the dCache was brought up. The suspicion is that these were lost files in the DPM era and ATLAS has not picked up on this until now, but I'm going to need to investigate this.
https://ggus.eu/index.php?mode=ticket_info&ticket_id=160308
 - if you checked that namespace import log file doesn't contain any warning than files were already lost before migration (e.g. consistency checks – migration first step – remove from DPM namespace files with no physical replica on disknode => this makes already lost files more visible)

- IR-IPM-HEP (ROC_Asia/Pacific)
 1. internet interruptions on irregular basis due to political reasons in Iran.
 2. they asked to postpone the deadline to June/23.



Thank you

Questions?

www.egi.eu



This work is partially funded by the EU research and innovation programme