

Migration DPM → EOS, dCache retour point de vue opérationnel ATLAS

Frédéric Derue, LPNHE Paris

Journées LCG-France, LPNHE Paris
7 Juin 2023

An ATLAS DDM requirement

« For storage migrations, please open a JIRA ticket in ADCINFR »

● **DPM → EOS migration**

- concerned GRIF
- not only a migration but also creation of a single Storage Element
- JIRA tickets <https://its.cern.ch/jira/browse/ATLDDMOPS-5603> and <https://its.cern.ch/jira/browse/ADCINFR-236>
 - done since months, ~May-June 2022 for IJCLab/LPNHE, September for IRFU but still ongoing for the other experiments (CMS, CTA ...)

● **DPM → dCache migration**

<https://twiki.cern.ch/twiki/bin/view/DPM/DpmDCache>

- CPPM : JIRA tickets <https://its.cern.ch/jira/browse/ATLDDMOPS-5641>
 - ticket opened on 8th March, closed on 28th March
- LAPP : JIRA tickets <https://its.cern.ch/jira/browse/ATLDDMOPS-5636>
 - ticket opened on 1st Feb, then long period before shutdown mid-March transfers are ongoing
 - there existed 2 SEs at same time for a while: one for DPM, one for dCache
 - ⇒ only one now for dCache
 - dumps for the localgroupdisk area is waiting, then ticket can be closes
- LPC : no JIRA tickets, by GGUS ticket ?
 - Using wiki/instructions from Petr
 - done ~20th March

● DPM → dCache migration

- Tokyo : JIRA ticket <https://its.cern.ch/jira/browse/ATLDDMOPS-5639>
https://ggus.eu/index.php?mode=ticket_info&ticket_id=158821#update#11
 Done in March 2023 after a short downtime
- RO-07 : JIRA ticket <https://its.cern.ch/jira/browse/ATLDDMOPS-5607>
 Done end of May-early June 2022 after a short downtime
- Beijing : https://ggus.eu/index.php?mode=ticket_info&ticket_id=158787
 ticket opened in 09/2022, migration done in May 2023 during a downtime
 no JIRA ticket opened, several private emails
 Panda Queues do not seem to have restarted
 ⇒ total absence of information exchange with FR-cloud
- Hong-Kong ?

In general most/all work done centrally by DDM-ops – and ADC –
 concerning instructions and setup of
 Storage Elements, Computing Elements, Squid, to be done in CRIC
 Minimal work/follow-up from FR-cloud support

CPPM


157817
transfers failures as source
⇒ files to be declared as lost

160998 25/03-05/04
Low transfer efficiency
⇒ long thread, full disk due to
removal of dpm names (rm -f)
dumps of disks provided

161661 25/04
squid down
⇒ site was in
downtime

161923 12-15th May
large spike in errors
between IN2P3-CPPM
(source) - BNL (dest)
⇒ « modification to allow
a new vo had
impredicted effects »

Slots of Running jobs (HS06)



158454
ATLAS data access pb (EventIndex)
⇒ two disk server had xrootd daemon problems


160290
failed to transfer files
⇒ downgrade dpm and xroot
solve the issue

DMP → dCache

List of **GGUS ticket** to IN2P3-CPPM related to ATLAS

CPPM

DMP → dCache



160998 25/03-05/04

161661 25/04

161923 12-15th May

LAPP grid CPU (last 12 months)

LAPP

160956 05/04
deletion errors
⇒ DPM server restarted


160985 23/03-01/05
161028 11/04
161474 20/04
squid issues

161591 18-20/04
transfer/deletion
errors
⇒ more servers
added for
dCache services

161865 7-11th May
transfer/deletion errors
⇒ disk servers restarted
+ DPM certificate
renewals

157794
transfer failures
⇒ certificate renewal

160785 05/03
transfer/deletion errors
⇒ hardware issues



158548 29/08
transfer failures
⇒ IPV6 pb + new disk servers

X-mas scheduled downtime

159927 15/12
⇒ storage « at risk »
during Xmas


spring scheduled downtime

DMP → dCache
2 SEs at same time

List of **GGUS ticket** to IN2P3-LAPP related to ATLAS

LAPP

Start of
DMP → dCache



160785 05/03

160985 23/03-01/05
161028 11/04
161474 20/04
squid issues

161591 18-20/04

161865 7-11th May
transfer/deletion errors
⇒ disk servers restarted
+ DPM certificate renewals

LPC

157947 04-25/07
high failure rate of
production jobs,
LRMS error:
⇒ restart of services

158985 23/09-14/11
killed jobs by condor

159562 20/11-02/01
new ARC-CE

160527 14-21/02
deletion errors

160715 27/02-01/03
transfer / deletion
⇒ reboot

161004 26/03-
30/03
transfer/deletion
errors
⇒ post-migration

161668 25/04-02/05
transfer/deletion errors
⇒ change ACLs

Slots of Running jobs (HS06)

DMP → dCache

158382 04/08
transfer/deletion
⇒ new disk pool ?

158442 08/08
squid issues
⇒ new hardware

158853 28/08-13/09
high job failures due
to "LRMS error"

List of **GGUS ticket** to IN2P3-LPC related to ATLAS

LPC

DMP → dCache




161004 26/03-30/03
transfer/deletion
errors
⇒ post-migration

161668 25/04-02/05
transfer/deletion errors
⇒ change ACLs

GRIF-IRFU grid CPU (last 12 months)

GRIF-IRFU



158059 17/07-24/09 transfer errors
=> during migration

DMP → EOS


161971 17-18/05 transfer errors
=> power-outage

161961 17-30/05 squid down
=> power-outage


List of **GGUS ticket** to GRIF related to ATLAS

GRIF-IRFU

DMP → EOS in summer 2022



GRIF-LAL



List of **GGUS ticket** to GRIF related to ATLAS

GRIF-LAL

DMP → EOS in summer 2022



Why so low efficiencies ? Not related to any GGUS ticket

GRIF-LPNHE

157947 04-25/07
high failure rate of
production jobs,
LRMS error:
⇒ restart of services

159336 26/10
transfer low
⇒ glitch in
auth/auth/z from the
side of FTS3 and
TPC workflow

159357 29/10-06/11
transfer errors
⇒ bad data balance
amongst the EOS
filesystem and some
filesystem was
100% full

159846 08-12/12
transfer errors
⇒ no answer

160706 25/02
transfer errors
⇒ hi-load on
eos-mgm

161634 23/04
transfer errors
⇒ glitch ...

162203 05/06
Files inaccessible

160993 24/03
transfer errors
⇒ no answer

159988 23/12-06/02
transfer errors
⇒ allocated space and
quota disagreement
⇒ eos upgrade

158872 12/09
transfer low
⇒ missing DN

158778 08/09
transfer low
⇒ not clear if site pb

159512 16/11
transfer errors
⇒ issue with
quota level

Slots of Running jobs (HS06)

Pledge GRIF

Pledge GRIF-LPNHE~30% GRIF

DMP → EOS


157814 22/06
transfer failures
⇒ after rm of storage

158096 20/07
squid problems
⇒ cooling issues

List of **GGUS ticket** to GRIF related to ATLAS


GRIF-LPNHE

DMP → EOS in summer 2022




Since 1 year

[link](#)




Since March

[link](#)



T2-France

- Pledges
- IN2P3-LAPP
- IN2P3-CPPM
- GRIF-LPNHE
- GRIF-IRFU
- IN2P3-LPC
- GRIF-LAL



- **Pledge for France (see [cric](#))**

- in 2023 increases by 4%, represents 9.4% of all T2s

- **LOCALGROUP** [live-storage](#)

- non pledged resources

	Total (TB)	Free (TB)
CPPM	245	187
GRIF	700	322
LAPP	58	15
LPC	28	10
Beijing	0	0
Tokyo	2000	914
RO-07	11	3

[live-storage](#)

large part not yet reinstalled

	Pledge 2022 (TB)	Pledge 2023 (TB)	Installed (TB)
CPPM	2200	2200	2264
GRIF	5510	6099	4150
LAPP	3700	4500	5812
LPC	2188	2188	2197
Beijing	400	400	0
Hong Kong	1050	1050	1056
Tokyo	8000	8800	10400
RO-07	3500	3500	3804

After migration DPM → dCache / EOS

- CPPM, LPC came back to their previous storage size
- LAPP increased by ~1.3 PB !
- GRIF still miss ~2 PB !

- **DPM → dCache : CPPM, LAPP, LPC**

- done around March in one go for CPPM, LPC with two co-existing SE (DPM and dCache) for LAPP
- not so easy to spot operational issues looking at « Status » or « cpu » plots
- CPPM cpu plot is more shaky in 2023 than in 2022 – but not related to migration
- LAPP storage gained unexpected 1.5 PB after migration
- number of GGUS tickets during last year
CPPM (6), LAPP (9), LPC (10)

- **DPM → EOS + SE unification : GRIF**

- done around May-June 2022 for IJCLab/LPNHE and July-September for IRFU
- ~20 GGUS tickets – for three labs – since last year
⇒ only Manoulis seem to answer to these tickets
- really hard to relate tickets to drops of cpu in plots – not really related ?
- IJCLab amount of cpu is strongly decreasing since March !
- GRIF ATLAS storage pledge still not in place whereas migration done one year ago ⇒ still ~2PB missing !

In general it is not easy to relate losses of efficiencies (cpu, data transfer ...) or (in particular) instabilities to GGUS tickets or known issues, in particular related to the migration