

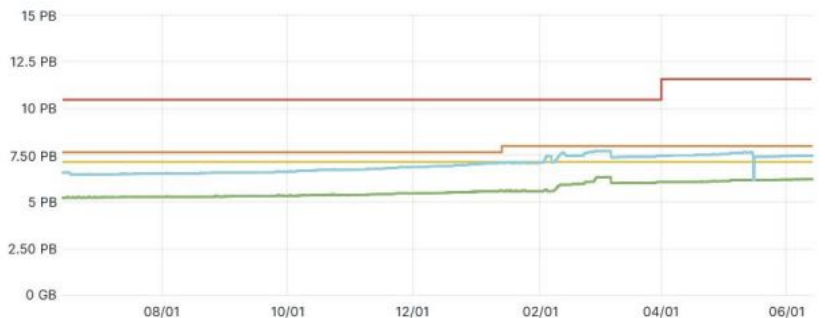
Report LHCb

Lucio Anderlini


actually, most contents were stolen
from Andrea Rendina (thanks!)

Business as usual (in no-data-taking conditions)

- Steady MC production (good job efficiency and pledge target achieved)
- Disk quota below pledge will not be an issue before October/November
- Well below tape pledges



	min	max	avg	current
disk used	1.400 PB	6.339 PB	4.605 PB	6.201 PB
disk quota	1.389 PB	7.133 PB	5.977 PB	7.133 PB
disk+buffer used	1.581 PB	7.739 PB	5.822 PB	7.467 PB
disk+buffer quota	1.691 MB	7.984 PB	3.835 PB	7.984 PB
pledge	10.461 PB	11.561 PB	10.683 PB	11.561 PB



used	11.928 PB
pledge	25.261 PB

LHCb visit at CNAF, yesterday

Many thanks to the **Tier-1 team** for welcoming Chris Haen (LHCb Data Manager), Concezio Bozzi (LHCb Computing Coordinator) and myself at CNAF, yesterday, for a review of the utilization of CNAF resources by LHCb.

Topics:

- Plans on the migration to Tecnopolo
- Status and perspective for farming (in particular non-x86 architectures)
- Development of Network
- Storage and Data Management

Open issue: re-organizing file system

- Put disk and buffer on different file systems
 - Buffer on SSD disk
 - How to define buffer size?
 - In the past buffer size was defined as the amount of data which LHCb could write in a few days
- In order to use the file protocol directly from the WNs, LHCb needs to use SRM
 - First discussed in https://ggus.eu/index.php?mode=ticket_info&ticket_id=157208
 - We could move

`/storage/gpfs_lhcb/lhcb/disk/LHCb` to
`/storage/gpfs_lhcb/lhcb/lhcb/LHCb`

during a short and scheduled downtime, unmounting the 3 filesets (`/storage/gpfs_lhcb/lhcb/[disk/user/tape]`), renaming them and mounting them again to the different mount points (to be agreed) in the same filesystem.

Future challenges

- **StoRM Tape REST API (StoRM-less tape recalls):**
 - Pre-production phase in our testbed
 - To be enabled for ATLAS production transfers soon, starting with LOCALGROUPTAPE (almost no transfers for this Rucio RSE) just to verify everything is configured correctly. Probably zero overlap with SRM and TAPE REST for production transfers.
- **Token-based AuthN/Z:**
 - Recently configured in StoRM WebDAV for both ATLAS and CMS following WLCG profile (scope-based authorization)
 - Recently configured in XrootD for CMS
(https://ggus.eu/index.php?mode=ticket_info&ticket_id=161981)

slide from *Andrea Rendina*¹²

Plan for a next Tape Challenge

Target: complete the challenge before the end of September

1. Remount `/storage/gpfs_lhcb/lhcb/disk/LHCb` → `/storage/gpfs_lhcb/lhcb/lhcb/LHCb`
scheduling downtime if needed
2. We run a validation test on a CNAF *user interface* to ensure everything is in place
3. We restart the production and validate
4. If in the meanwhile developments for Storm TAPE Rest API get ready, we include them also for LHCb
5. We run the tape challenge

Moving buffer to SSD would not cause any issue for LHCb, but might require more time CNAF side: test postponed

Token-based AuthN/Z for storage requires some additional discussion and agreement on the policy to adopt. Definitely not for this Tape Challenge. CNAF agreed on being a “test-bed”.