



# Storage Issues: the experiments' perspective

Flavia Donno  
CERN/IT

*WLCG Grid Deployment Board,  
CERN 9 September 2008*

# The questions

---

- What are the main storage issues that compromise data transfer or access ?
- What are the operational problems that imply loss of times and improper use of human time and resources ?



# ALICE

---

- Alice is using as storage system xrootd enabled storage for all possible systems: dcache, castor, dpm and pure xrootd storage.
- Good performance for all systems and very good operational procedures. No issues to report
- Alice would like to trigger the question of the policy for SE upgrades, taking into account the new developments in terms of xrootd: how are the upgrades planned and are they going to be rolled out for the different flavors of the SEs? Finally, is support for xrootd going to be guaranteed ?



# ATLAS

---

## ■ Configuration issues

- Difficult to understand how to address sites to configure specific resources for ATLAS and very difficult to find a common configuration (CASTOR/dCache)
- While the interaction with the developers and deployment team is efficient for CASTOR, it is less efficient for dCache
- For dCache the knowledge is sometimes in the hand of one developer who can be unavailable. A global view of the system behavior seems missing. Documentation (user guide and system administration) seems missing.
- A reference installation where to try things on is missing for dCache. CASTOR offers many tools to try possible configuration and the FIO assistance is good. dCache is weaker in this respect.

## ■ Certification and release procedures

- ATLAS suffered many times for sites being unavailable because of a destructive upgrade (more frequent in dCache than in CASTOR but more dCache sites as well). Need for stronger certification procedures
- WLCG high-level tools should be part of the certification procedure reports).

■ SRM MoU short term plan needs to be in place as programmed

■ Need to protect spaces, tape usage, etc.



# CMS

---

## ■ *Functionality*

- Most of the issues related to missing functionalities are due to the lack of proper mechanisms to handle storage quotas and file/disk ACLs. This is going to hit CMS as soon as the analysis activities will increase. Thinking to move toward the usage of space tokens

## ■ *Operations*

- Mainly related to the states in which the system may go as a consequence of operations not permitted by the CMS model but permitted by the storage systems.
  - CMS deployed different Phedex end-points for T0, T1 and CAF in order to address the needs of the data operations. Unfortunately sometimes (e.g. in data subscriptions from remote T1 sites to >1 PhEDEx nodes at CERN) Castor triggered pool-to-pool copies that causes errors like "file exists" at the CMS data management level.
  - CMS jobs access data directly on the storage systems to avoid a copy of the input files to the local WN disks. This means that the MSS may be contacted multiple times for the same file.



# LHCb

## ■ *Site configuration and stability*

- In general difficult. More problems with dCache than with CASTOR.
  - CASTOR at CNAF: Failed to perform gfal\_ls: [SE][Ls] httpg://srm-v2.cr.cnaf.infn.it:8443/srm/managerv2: Too many threads busy with Castor at the moment Unknown error 18446744073709551615
  - CASTOR at RAL: ORACLE problem.
  - NIKHEF: we had so many issue that are all documented by several GGUS tickets.
    - 1. PinManager issue [https://gus.fzk.de/pages/ticket\\_details.php?ticket=39549](https://gus.fzk.de/pages/ticket_details.php?ticket=39549),
    - 2. SRM dead: [https://gus.fzk.de/pages/ticket\\_details.php?ticket=39996](https://gus.fzk.de/pages/ticket_details.php?ticket=39996)
    - 3. (POLLIN) Control line creation failure [https://gus.fzk.de/ws/ticket\\_info.php?ticket=39953](https://gus.fzk.de/ws/ticket_info.php?ticket=39953)
    - 4. lcg-gt timing out (becasue LCG\_GFAL\_INFOSYS was pointing to local dismissed top bdii kraal.nikhef.nl:2170)
    - 5. File ONLINE with not tURLS: [https://gus.fzk.de/pages/ticket\\_details.php?ticket=40001](https://gus.fzk.de/pages/ticket_details.php?ticket=40001)
    - 6. gfal\_turlsfromsurls: [SE][StatusOfGetRequest] httpg://srm.grid.sara.nl:8443/srm/managerv2: org.xml.sax.SAXParseException: An invalid XML character (Unicode: 0x10) was found in the element content of the document. Unknown error 4294967295 **\*\*ALARM\*\*** ticket [https://gus.fzk.de/pages/ticket\\_details.php?ticket=40170](https://gus.fzk.de/pages/ticket_details.php?ticket=40170)
  - IN2P3:
    - 1. Disk files NEARLINE at IN2P3: files supposed to be permanent on disk at IN2P3 are only NEARLINE on the storage. (dcache server down)
    - 2. gsidcap fails using shipped version of dcache clients not containing the fix for the concurrent gsidcap connections
    - 3. Jobs fail to store files in the remote storage at IN2P3 .



# LHCb

---

- IN2P3:
  - 4. [https://gus.fzk.de/pages/ticket\\_details.php?ticket=36999](https://gus.fzk.de/pages/ticket_details.php?ticket=36999) , CGI-gSOAP:Error reading token data: Success \*\*\*
- GRIDKA :SAM tests failing [https://gus.fzk.de/pages/ticket\\_details.php?ticket=39551](https://gus.fzk.de/pages/ticket_details.php?ticket=39551)  
gfal\_ls:Error reading tokendata: Success issue:  
[https://gus.fzk.de/pages/ticket\\_details.php?ticket=39995](https://gus.fzk.de/pages/ticket_details.php?ticket=39995)
- PIC: gfal\_ls:Error reading tokendata: Success issue:  
<http://lblogbook.cern.ch/Operations/416> (no GGUS open because like for GridKA this is not reproducible by changing the client versions)



**Thank You**



*WLCG Grid Deployment Board,  
CERN 9 July 2008*