

# DPM deployment

---

*M. Serra – INFN Roma*



## Outline: DPM deployment

---


- Since SRM 2.2 is not in production yet I will shortly summarize current DPM deployment status
  - in depth technical discussion in the BOF DPM afternoon session
- Some storage statistics from information system
  - top BDII(s): [lcg-bdii.cern.ch](http://lcg-bdii.cern.ch), [egee-bdii.cnaf.infn.it](http://egee-bdii.cnaf.infn.it)
- Some info collected via a questionnaire hand out to WLCG administrators
  - only few answers ☹ (~15)



## How many DPM around?

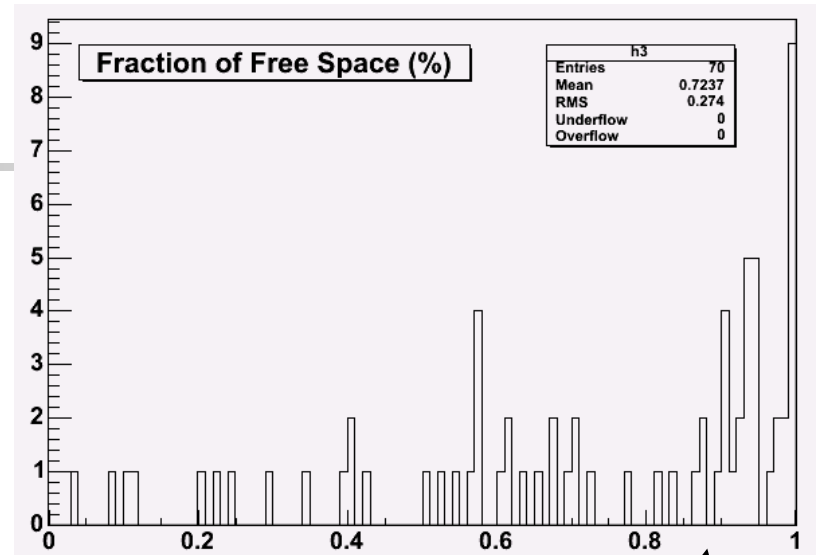
---

- Info collected from what is published in the information system (07-01-18)
  - NOT easy to understand if a storage is a DPM (or another type), no clear GLUE schema field available
- WLCG: ~81 DPM storage / ~305 storage system  
(including "classic SE")  
~65% of the "srm-SE"s , ~27% total
  - 11 in UK, 10 in Italy, 8 in France, 6 in Taiwan, 5 Greece...  
DPMs in all LCG "states" but none in US
  - 11/81 of them publish 0 total space and 0 free space and  
at least 6 of them were not down for maintenance
    - 1/81 ~1750 PB total space ..... wow!

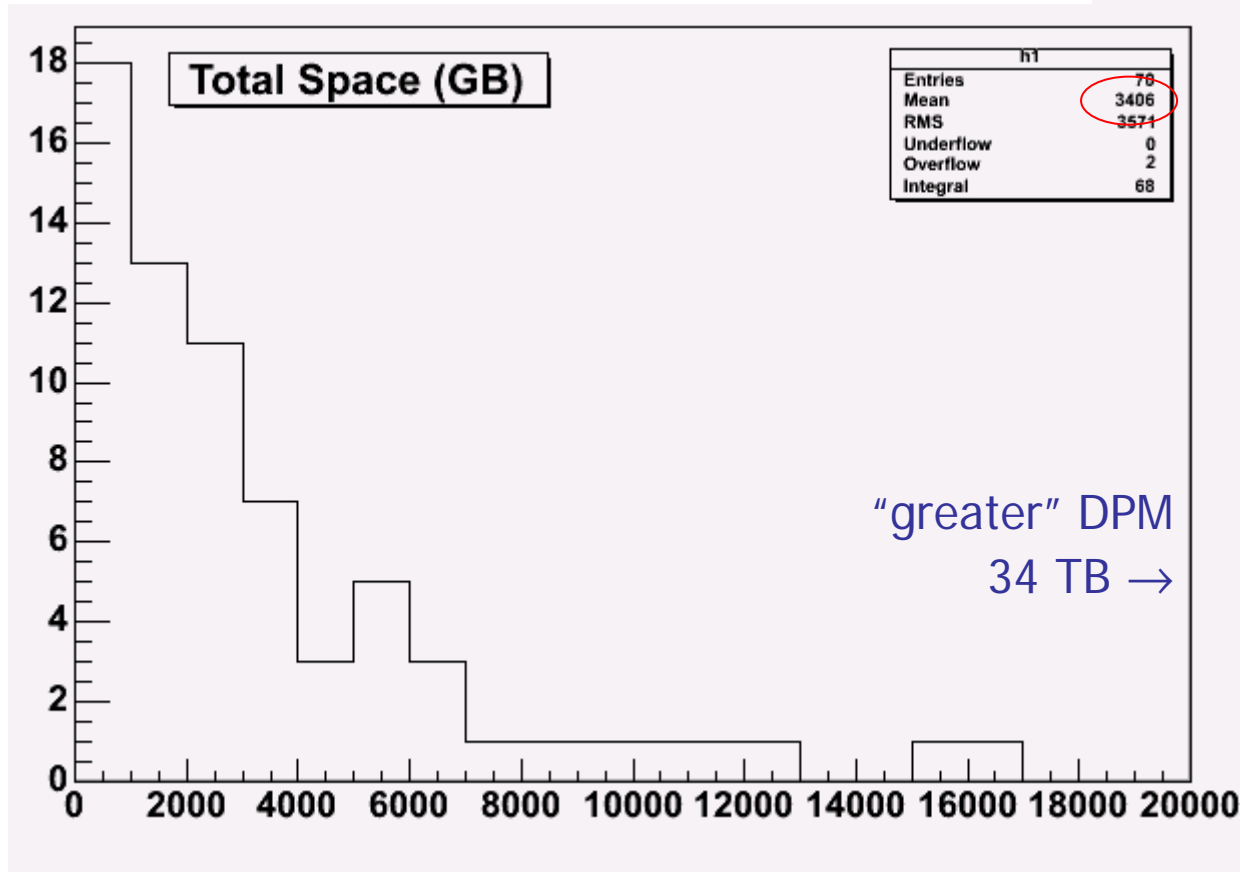


# DPM disk size

- Few installations > 10 TB



are people using these disks?





## Results from DPM survey – wishes (1)

---

- Simple tool to assess if DPM is “consistent”
  - check + correct, LFN  $\leftrightarrow$  PFN mapping
- Monitoring
  - real time & statistic, per VO, per pool, ...
- More options for load balancing
  - more advanced balancing (not only round robin)
  - (automatic) replication of hot files
  - addressing different pools according to VO role
- Quotas
  - VO-based quota



## Results from DPM survey – wishes (2)

---

- OS: 64 bit, SLC4, ... 1 Windows (!) & 2 ZFS for pool disk ☺
- security: ratelimit to avoid DOS, .....
- Xrootd for Alice sites
- long term: migration tool to MSS



## bugs & problems (survey + experiments)

---

- 1 CMS site: db of ~700MB (~10<sup>6</sup> rows) slow down DPM
  - missing cleaning of old get/put request
  - problem temporary patched creating db indexes (thanks DPM developers 😊)
  - ⇒ do we have “significant” DPM installations to test scalability?
    - @ T2 we will have 50-300 TB disk pools!
- Atlas: in DPM no secondary groups → users and prod-users different write permission: prod-users *ro* permission on “normal” user dir
  - patch under testing/deployment (temporary workaround)
  - Atlas: support for secondary groups is seen as a necessity
  - ⇒ do we need more “advice” to deploy/configure DPM @ T2?
- “neverending story” of DPM & Castor rfiio incompatibility
  - is/was it really a DPM problem or a coordination problem?
  - the proper fix is not integrated in the release yet.....



# Conclusions (survey + experiments)

---

- DPM PROs

- easy to install, stable, no major problems until now, good support, VOMS support, in case of special disaster recovery it is not impossible to “hack” for sys-admin

- DPM CONs

- lack of advanced features & tools (see wishes)


- Issues

- are we testing stability & scalability at right/needed scale?
- will advanced features come in time for T2?
- ~2 FTE (?) for development & support are enough?

- All admins asked for more documentation & training

- help to have a better configuration, db recovery procedure, ...





# Reminder

do not miss the discussion about DPM in the BOF afternoon session

