


GridPP
UK Computing for Particle Physics

Storage Management

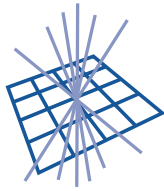
Greig A. Cowan

University of Edinburgh





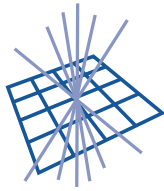
1. SRM v2.2 update
2. Storage monitoring with SAM
3. Storage accounting
4. DPM ACL change



- Directory functions, space reservation, support for access latency and retention policy.
- **Storage Classes** (T0D1, T1D1, T1D0).
 - VOs will use **space token** descriptions to identify classes (i.e. ATLAS_AOD).
 - GLUE 1.3 schema needed to advertise.
- CASTOR, dCache, DPM, StoRM... all implementing subset of spec.
- Functionality testing has been ongoing for months.
 - Stability has improved, now moving into production.

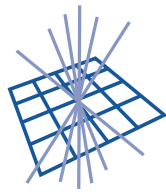
- dCache **v1.7.0-31** released this month.
 - SRM web service runs in Tomcat (needs Java 1.5).
 - SRM 1 and 2.2 endpoints available via same port (8443).
 - Full 2.2 functionality in v1.8 (April)
- DPM **v1.6.3** released this month.
 - Two different daemons: v1 on 8443 and v2.2 on 8446.
 - VOMs support.

http://goc.grid.sinica.edu.tw/gocwiki/How_to_publish_the_version_of_my_SRM



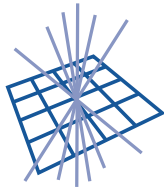
GridPP
UK Computing for Particle Physics

SAM monitoring of storage

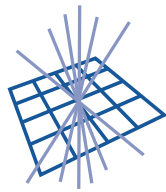


Existing SAM tests

- **Replica management** suite (CE-sft-lcg-rm).
- lcg_utils used to gridftp files to/from SE.

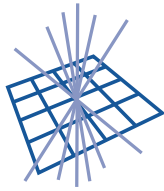


- Use case for VO **analysis** jobs.
- Open a file, read a few bytes. Open another file, read a few bytes. . .
 - GridFTP'ing set of files to WN is inefficient.
- DPM has `rfio`; dCache has `(gsi)dcap`. `xroot` also an option.
- **GFAL** provides common interface to these protocols.
- New SAM test (`CE-sft-posix`) to probe this access method.



Yesterdays results

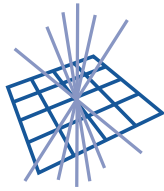
ScotGRID-Edinburgh	Pass	UKI-NORTHGRID-LANCS-HEP	Fail
UKI-SCOTGRID-DURHAM	Pass	UKI-NORTHGRID-LIV-HEP	Fail
UKI-SCOTGRID-GLASGOW	Pass	UKI-NORTHGRID-MAN-HEP	Fail
<hr/>			
UKI-LT2-Brunel	Pass	UKI-NORTHGRID-SHEF-HEP	Pass
<hr/>			
UKI-LT2-IC-HEP	Fail	UKI-SOUTHGRID-BHAM-HEP	Pass
UKI-LT2-IC-LESC	Pass	UKI-SOUTHGRID-BRIS-HEP	Pass
UKI-LT2-QMUL	Pass	UKI-SOUTHGRID-CAM-HEP	Pass
UKI-LT2-RHUL	Pass	UKI-SOUTHGRID-OX-HEP	Pass
UKI-LT2-UCL-CENTRAL	Pass	UKI-SOUTHGRID-RALPP	Pass
<hr/>			
UKI-LT2-UCL-HEP	Fail	RAL-LCG2	Fail



Problems observed

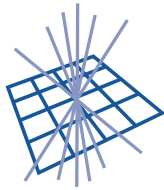
NorthGrid sites	Timeout during file read
UKI-LT2-IC-HEP	Transport endpoint not connected
UKI-LT2-UCL-HEP	No route to host
RAL-LCG2	SAM test has been changed to deal with CASTOR
UKI-SOUTHGRID-BHAM-HEP	No space left on device (intermittent)

- Will work with sites to resolve problems.

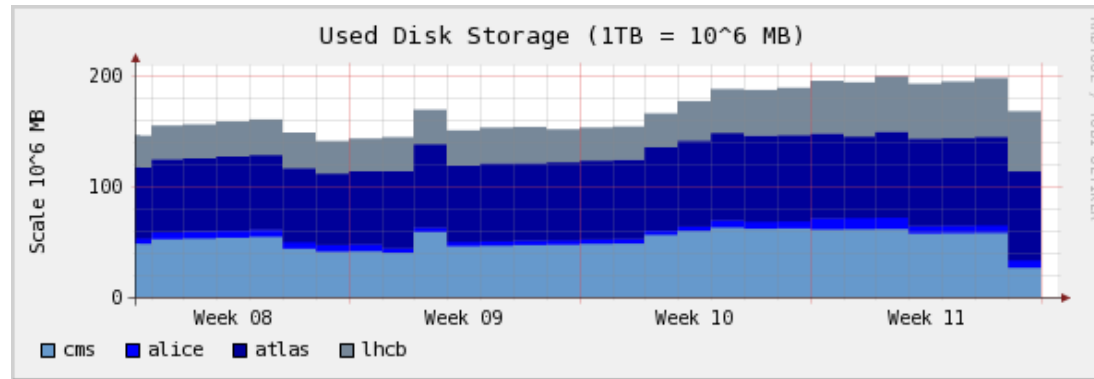


GridPP
UK Computing for Particle Physics

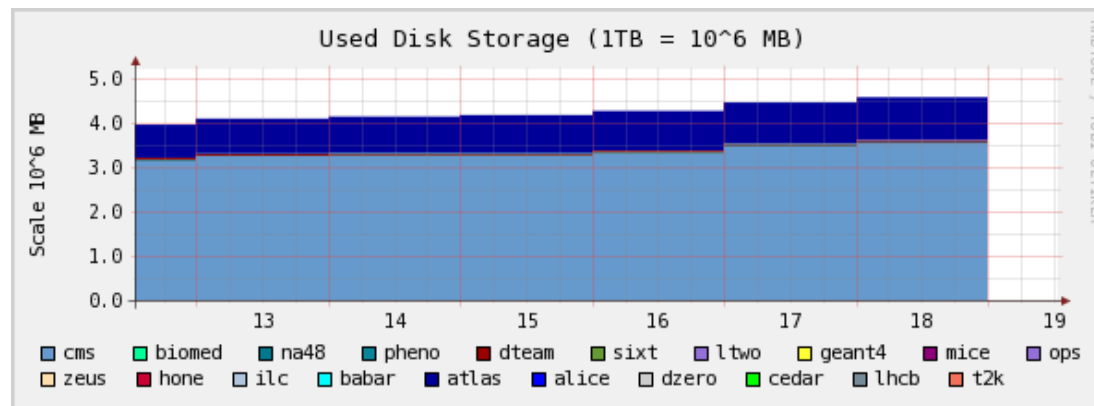
Storage accounting



GridPP disk usage over past month for LHC VOs ($>150\text{TB}$)



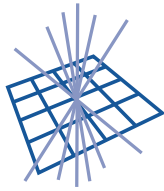
RHUL disk usage over past week for all VOs ($\sim 4.5\text{TB}$)



<http://goc02.grid-support.ac.uk/storage-accounting>

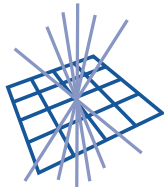
- Python script queries global BDII for `GlueSAStateUsedSpace`, `GlueSAStateAvailableSpace...` per site/VO.
- With SRM 2.2 the VOs will have storage classes.
 - Potential to collect accounting information at **space token** description level.
- Even with VOMs, user level information will not be available in BDII.
- Sites should check system for **consistency** - RALPP found problem.
- Open issues listed here:

http://www.gridpp.ac.uk/wiki/Storage_Accounting



GridPP
UK Computing for Particle Physics

DPM ACL change

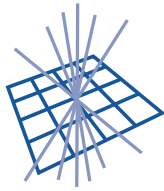


29/01/07 - First EGEE-BROADCAST

- All Atlas sites required to run the binary (no source code provided).
- Download from CERN AFS.
- Operation is “****EXTREMELY DELICATE****”. Bugs already found in Italy.
- I recommend that sites postpone until further notice. Contacted DPM developers.

06/02/07 - Binary run at Glasgow

- New bug found. Ordinary users could now not write to atlas/generated area.
- Graeme proposed fix.
- Working with sites to fix problem by hand.



08/02/07 - Atlas change their mind

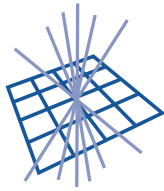
- Now only the atlas/dq2 directory tree to be changed.

15/02/07 - Second EGEE-BROADCAST

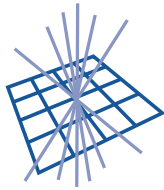
- New version of binary released.
- Now downloadable from voatlas01.cern.ch and AFS.

08/03/07 - Binary has ran at all GridPP sites.

- Still debugging problem at UCL-HEP.



- Interesting discussion in GGUS #18848.
- Middleware tools were insufficient for making change.
 - CLI `dpns-setacl` not recursive.
 - Faster to perform direct DB manipulation \Rightarrow binary was created.
- Are sites responsible for updating the ACLs for all VOs that they support?
 - Secondary VOMs groups should help.
- Again, this raises questions regarding site-VO contacts/communication.



- SRM 2.2 starting to be deployed.
 - Still need information from VOs regarding the space token descriptions.
- New posix SAM test.
- Storage accounting. Report inconsistencies to me.
- DPM ACL change raised new storage management questions.
- Latest information in the wiki, email list, meetings and the blog:

<http://gridpp-storage.blogspot.com>