GridPP
UK Computing for Particle Physics

Storage Accounting - update

Greig Cowan

University of Edinburgh

GridPP
UK Computing for Particle Physics

Outline

1. Background
2. The accounting system and front-end
3. Over-counting the space
4. Future considerations
5. Summary

GridPP
UK Computing for Particle Physics

Background

- Extensive levels of CPU accounting available. What about storage?
- Different user communities have different questions:
  - **Which sites are meeting their MoU targets?**
  - **Which SE's have space available for my production data?**
  - **Which VOs (and VO sub-groups) are using the storage at my site?**
- GridPP started prototype system for measurement/accounting.
  - Dave Kant and myself.

Details

- Every SE runs a **generic information provider** (GIP) which publishes information according to GLUE schema (currently v1.2).
- Concept of storage areas (typically 1 SA per VO).

```
GlueSiteName GlueSEArchitecture GlueSAStateUsedSpace
GlueSAStateAvailableSpace GlueSAType GlueSAPath
```

- Harvest these attributes by querying a top level BDII $\longrightarrow$ MySQL DB.
- This is a lightweight system.
  - **Nothing extra has to be installed on the SE.**

Web-portal

http://goc02.grid-support.ac.uk/storage-accounting/view.php

- Allows site admins/production managers/ROC managers to query the DB.
- Dynamically generates **historical plots** of the used space (**see later**).
- Shows table of latest used and allocated space (**see later**).
- **(Allocated space for VO) = (VO used space) + (free space available to VO)**

GridPP
UK Computing for Particle Physics

Front-end

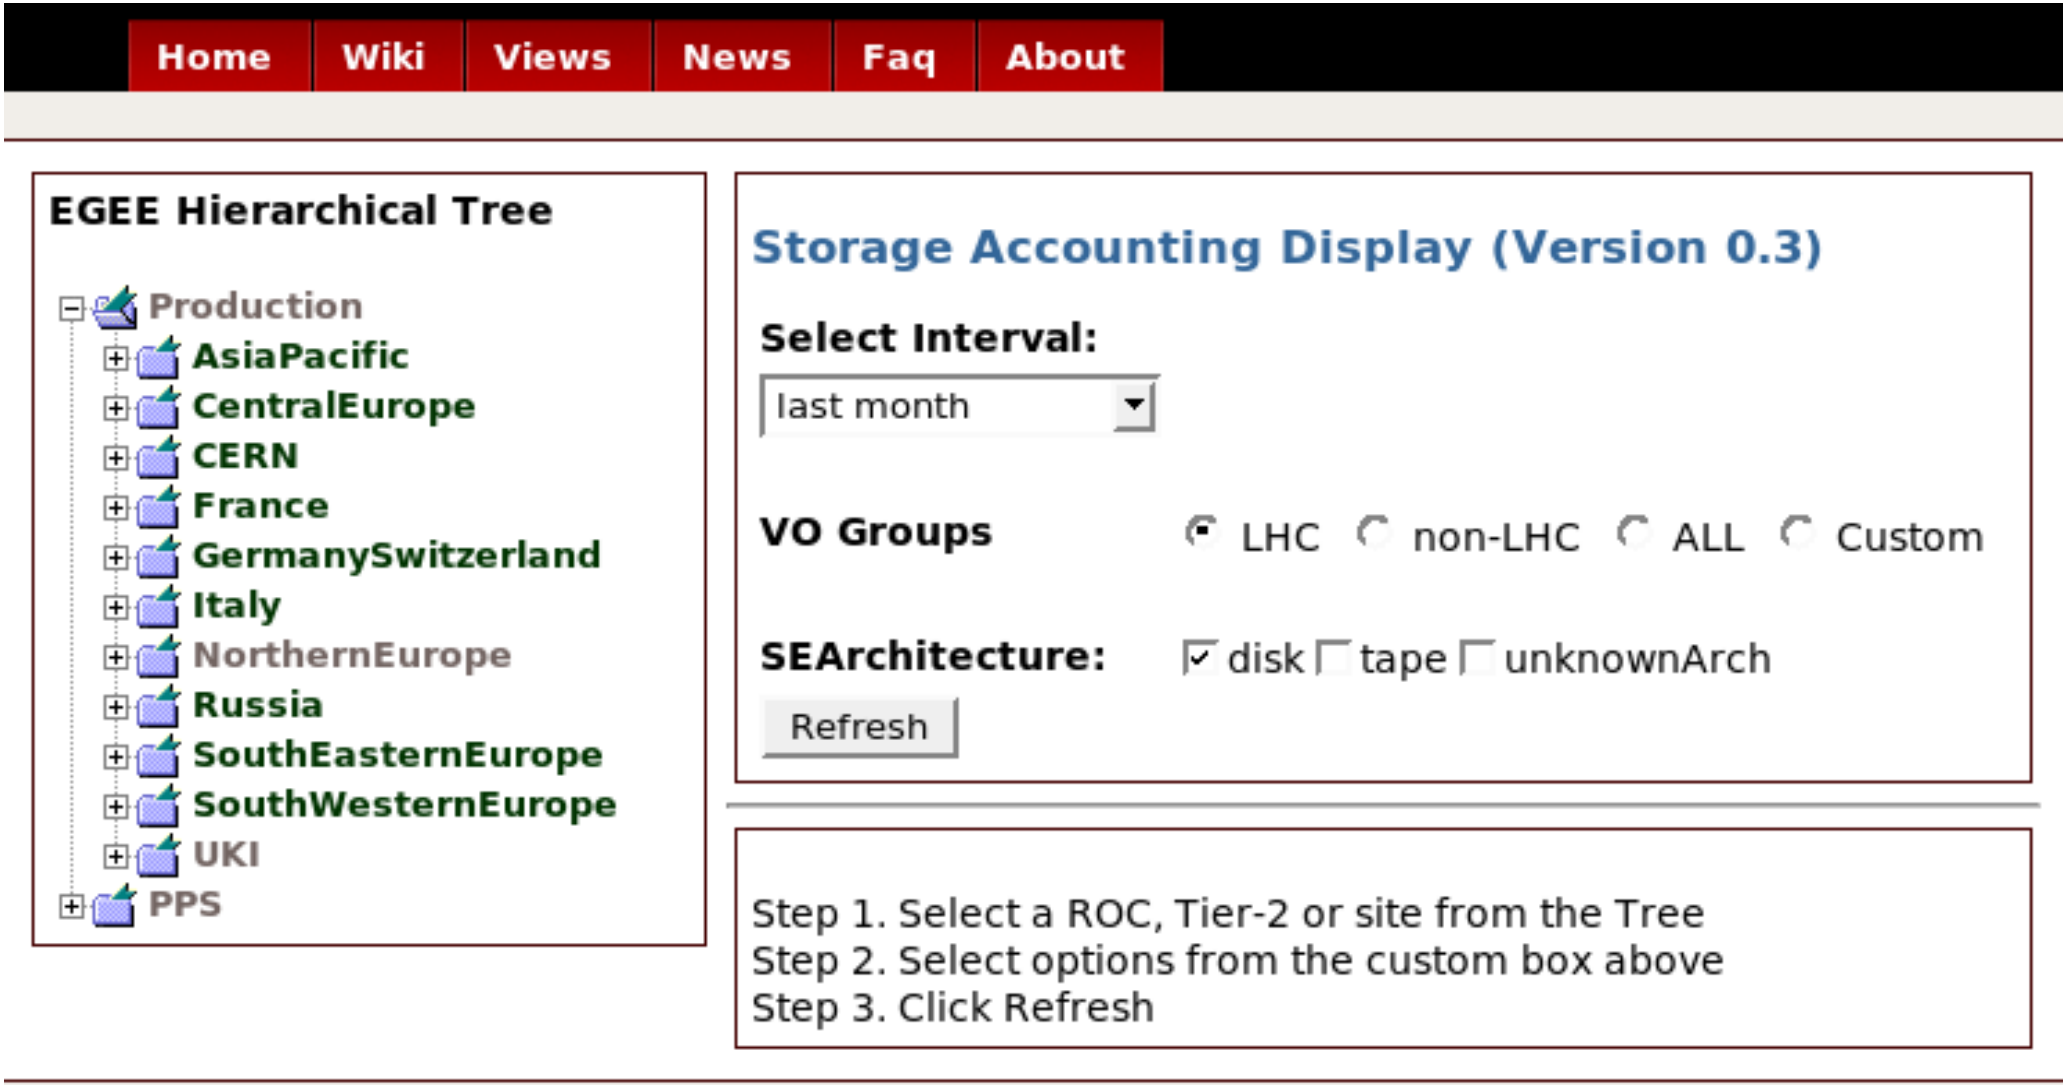

EGEE tree

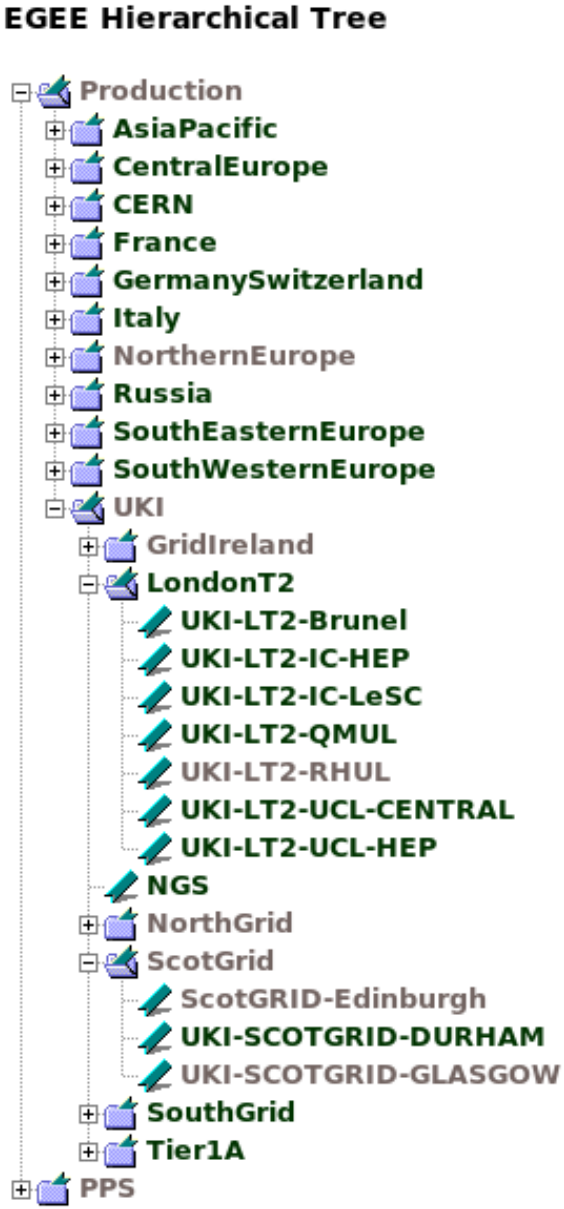
EGEE Hierarchical Tree
Production
AsiaPacific
CentralEurope
CERN
France
GermanySwitzerland
Italy
NorthernEurope
Russia
SouthEasternEurope
SouthWesternEurope
UKI
GridIreland
LondonT2
UKI-LT2-Brunel
UKI-LT2-IC-HEP
UKI-LT2-IC-LeSC
UKI-LT2-QMUL
UKI-LT2-RHUL
UKI-LT2-UCL-CENTRAL
UKI-LT2-UCL-HEP
NGS
NorthGrid
ScotGrid
ScotGRID-Edinburgh
UKI-SCOTGRID-DURHAM
UKI-SCOTGRID-GLASGOW
SouthGrid
Tier1A
PPS

GridPP
UK Computing for Particle Physics

VO selection

Storage Accounting Display (Version 0.3)

Select Interval:

last month

VO Groups ○ LHC ○ non-LHC ○ ALL ○ Custom

VOs:

☐ alice	☐ atlas	☐ babar	☐ biomed	☐ cdf
☐ cedar	☐ cms	☐ cosmo	☐ dteam	☐ dzero
☐ egeode	☐ esr	☐ fusion	☐ geant4	☐ gear
☐ gene	☐ gin	☐ gitest	☐ gridpp	☐ hone
☐ ilc	☐ lhcb	☐ ltwo	☐ magic	☐ manmace
☐ mariachi	☐ marine	☐ mice	☐ minos	☐ na48
☐ ngs	☐ ops	☐ oxg	☐ pheno	☐ planck
☐ ralpp	☐ sixt	☐ solovo	☐ swetest	☐ t2k
☐ webcom	☐ zeus			

SEArchitecture: ☑ disk ☐ tape ☐ unknownArch

Refresh

Site/ROC historical usage

GridPP disk usage over past month for all VOs (~ 400TB)

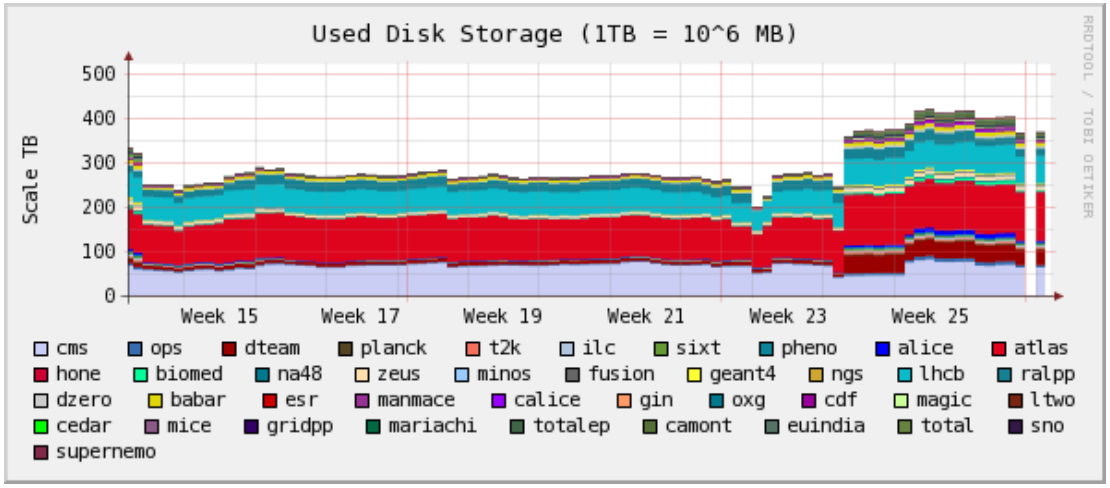

NGDF-T1 disk usage over past 3 months for LHC VOs (> 100TB)

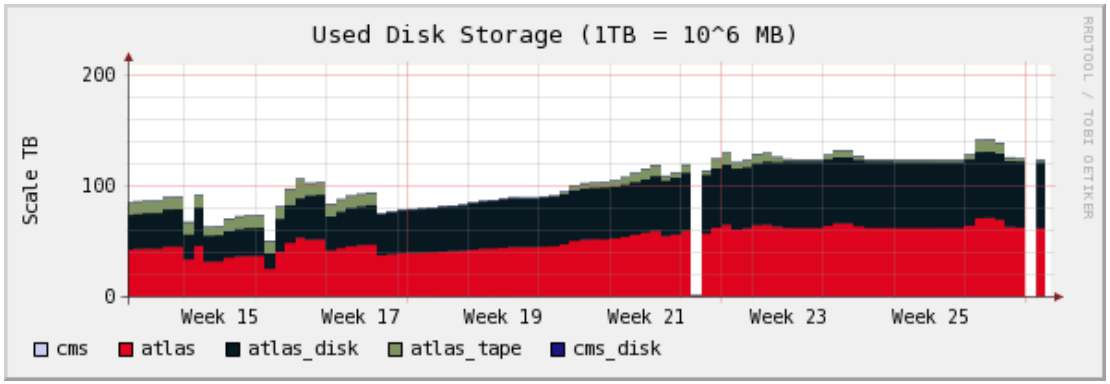

VO historical usage

LHCb disk usage over past 3 months in GridPP (> 60TB)

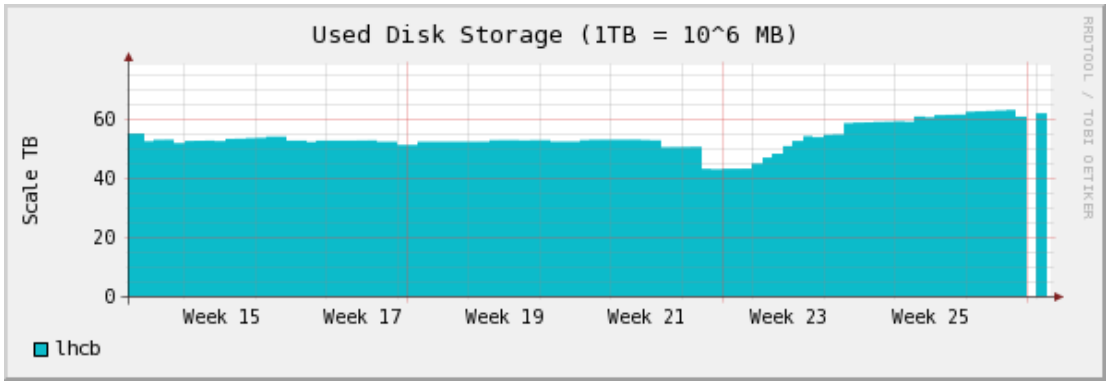

ATLAS disk usage over past 3 months in GridPP (> 100TB)

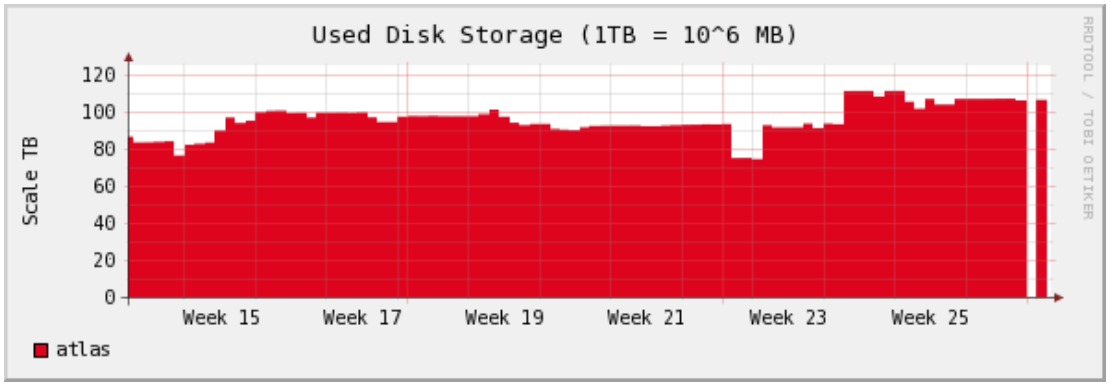

GridPP
UK Computing for Particle Physics

Project organisation

- Implementation details:

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

- Bug tracker:

  https://savannah.cern.ch/projects/storage-account/

Over-counting of storage

- If VOs **share** DPM pools or dCache pool groups then the default GIPs do **not** correctly report the available and used space per VO.

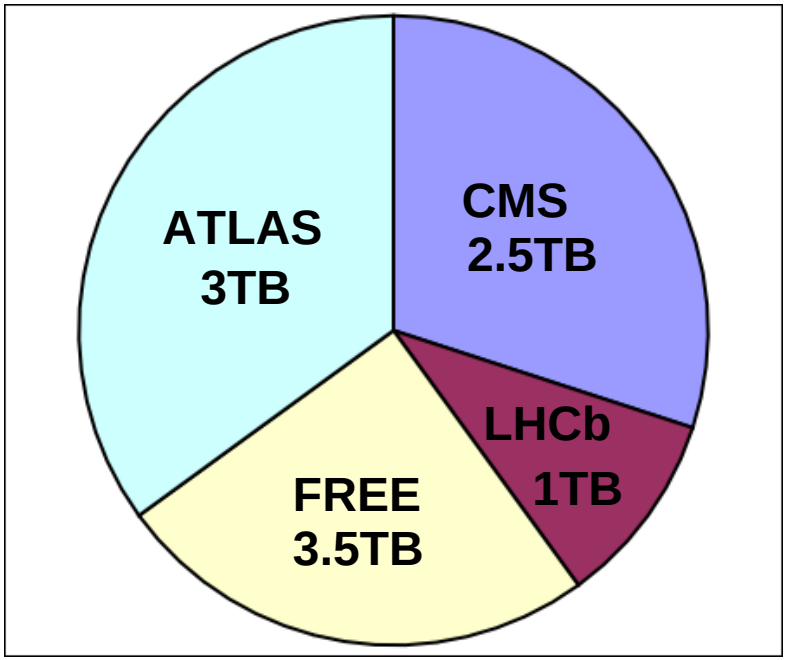

For dCache:

- All VOs reported as having **6.5TB used**.
- All VOs reported as having **3.5TB available/free**.
- VOs must be assigned their own pools/pool groups.

Over-counting of storage

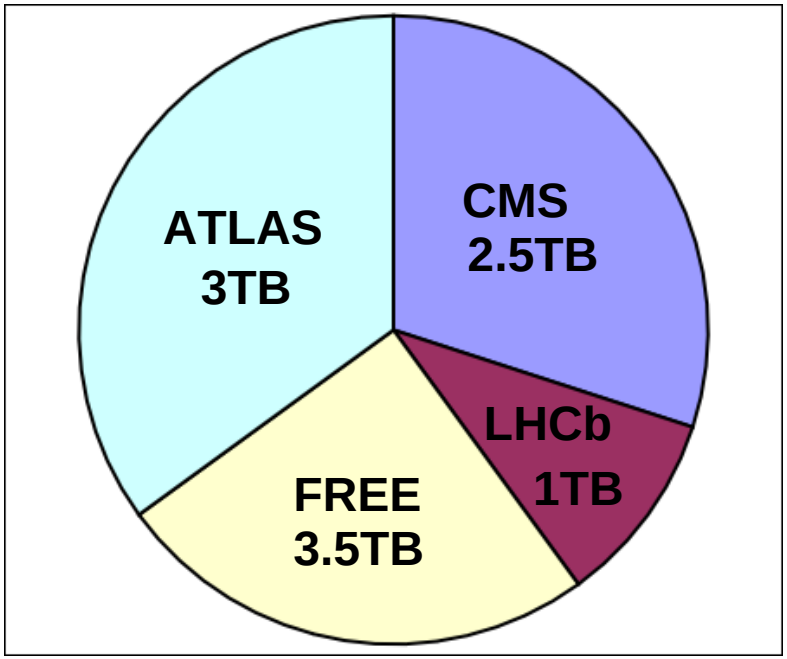

For DPM:

- $>$ v1.6.4 comes with an improved plugin.
- **Correctly** reports used space per VO.
  - performs DB query.
- All VOs reported as having **3.5TB available/free**.

- Reporting of the “same” free space is not necessarily a problem.
  - Users of system need to be aware of the situation.

DPM GIP plugin

UKI-SCOTGRID-GLA

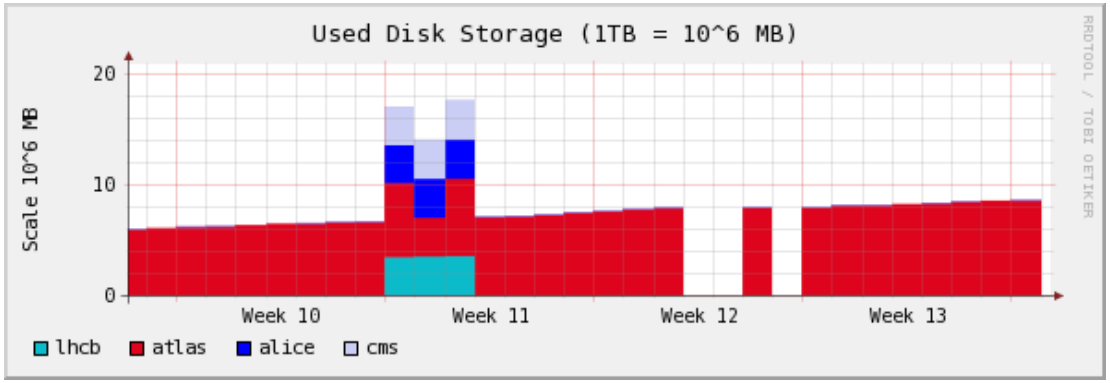

- Notice the blip when new plugin turned off.
- **Sites should check that information being correctly reported**.
- DPM developers already started work on internal accounting/quotas.

GridPP
UK Computing for Particle Physics

Input from experiments

- ATLAS and LHCb have similar needs. Want to know:
  1. **which sites have their existing (production) data.**
  2. **which SEs have free space for them to use.**
- Not concerned too much with detailed accounting (at the moment).
  - Maybe we want user-level storage accounting at some stage?
- Additional comments welcome.

GridPP
UK Computing for Particle Physics

SRM 2.2 and GLUE 1.3

- With SRM2.2, will be able to reserve spaces on SRMs.
  - **Static or dynamic depending on implementation.**
  - **Potential within GLUE 1.3 for each SRM to advertise space token descriptions.**
    * e.g, `ATLAS_AOD`
  - **Should be possible to account for storage at this level.**
- Sites need documentation to configure spaces.
  - **GSSD collecting this from experiments and developers.**
- GIPs need to be written for all of the SRMs.

GridPP
UK Computing for Particle Physics

Future plans

- Additional views of the data to be added, e.g.,
  - **ROC level will show contributing sites, not individual VOs.**
  - **VO-view for selected ROCs.**
- Support the move to GLUE schema v1.3.
- Do we want/need user-level accounting?
  - **Information system was designed for resource discovery, not accounting.**
  - **IS will not scale to publishing for each user.**
  - **If yes, then maybe the sensor should ask the SRM and not BDII for information.**
    - **Extension to SRM spec???**