

Data Access for Analysis

*Jeff Templon
PDP Groep, NIKHEF*

*A. Tsaregorodtsev, F. Carminati, D. Liko,
R. Trompert*

GDB Meeting
8 feb 2006

Use Case: Reprocessing or Analysis

- ◆ Job sent to site whose SE has (or is "close"?) to data
- ◆ Job shows up and needs to "do something" prior to "open"
- ◆ What???

Approaches seen in the wild

- ◆ “copy to local storage” : lcg-cp to TMPDIR or equivalent, read from there (nb good to know where local storage is!)
- ◆ (gsi)rfio : apparently supported by all DPMs but spotty elsewhere
- ◆ (gsi)dcap: horribly insecure
- ◆ Xrootd: only used by ALICE?
- ◆ GFAL: seems to be somewhat unknown in app community

HEPCAL (2002)

The issues regarding DS access in support of analysis jobs are largely addressed in HEPCAL, which assumed that the Data Management System would transparently optimize data access on the user's behalf. HEPCAL anticipated that at least the following options would be considered by the DMS:

1. Access (possibly via remote protocol) to an existing physical copy of the DS;
2. Making a new replica to an SE – because this SE has file-systems mountable from the chosen worker node, or perhaps it supports the protocol requested by the application – and arranging for the user program to access this new one;
3. Making a local copy to temporary storage at the worker node where the job is running;
4. If a virtual definition of the dataset exists, materializing the DS to either a suitable SE or local temporary storage at the node where the job will run.

The user will in general not be aware of this; her program will just “open the DS”. Subsequent reads on the returned handle will “get the bytes”.

Seems the answer is GFAL

- ◆ Can we get a GFAL presentation soon in the GDB about
 - What it can do
 - What underlying site protocols (gsiftp / rfio / dcap / xrootd) it wraps
 - Vision for the future

Laurence's 3 issues + JT's 1

- ◆ If you put your file in this SA, what will be the 'response time' when you try to read it later?
- ◆ If you put your file in this SA, what is the guaranteed lifetime of this file?
- ◆ Online Storage Areas (disk 'buffer' fronting tape)
 - How to model?
 - How to find out what is 'on disk' and what isn't?
- ◆ Update 'conventions' in IS (invented in EDG) to 1.2, and update software to use the new info (vo:/path/to/area)

GLUE Schema Evolution

- ◆ Previous versions designed *in vacuo* -- no use cases
- ◆ Users:
 - Sites
 - Grid users
 - Middleware people
- ◆ If something cannot be done (*e.g.* a certain query): define a concrete use case
 - Drive schema (or use case) evolution
 - Use case provides validation check for *GSTAT* (does each *SE* in Grid provide sensible answer for this use case's query??)

Couple Obvious Things Missing

- ◆ Quality of Service information for SEs / SAs
 - File "safety": are they backed up? Mirrored? RAID? Right now last is in architecture as "multidisk" but probably not appropriate as arch, more as QoS
 - Number of available transfer (gridftp, dcap, etc) connections
 - What else??
- ◆ GSTAT checks on published information: 'regression tests' for observed errors

Related Problem (but not in IS)

- ◆ When I ask an SRM for a file, how do I find out how long it will take before I can open?? Relevant for HSM systems with OSAs.

Status of Task

- ◆ Document in progress, 'modeling storage resources with GLUE'
 - Addresses the exp't use cases discussed last time
 - Guide to publishing the GLUE information
- ◆ Add'l document on "GLUE discrepancies" (input for GSTAT monitor)
- ◆ Bugs being submitted on DM tools where they use info incorrectly
- ◆ Hope to get significant improvements for gLite 3.0