

Storage and the Grid


Oxford Joint GridPP NGS day

Jens Jensen

j.jensen@rl.ac.uk



Science & Technology
Facilities Council



GridPP

UK Computing for Particle Physics



Background

- SRM is the Storage Resource Manager
 - A *control* protocol for controlling mass storage
 - OGF standard (GFD.129)
- SE is a Storage Element
 - Grid interface to storage
 - Interfaces for *control, data transfer, and information*
 - By abuse of notation often called “an SRM”



Background

- SRB is the Storage Resource Broker from SDSC
 - A “data grid”
 - Commercial (from GA) or free-for-non-commercial
- iRODS wants to be Caliph instead of the Caliph
 - “New data grid technology”
 - Aiming for preservation and digital libraries
 - Rule based, running μ services



SRM (SEs with SRM interfaces)

PROS

- Open Standard (OGF)
- >6 interoperating versions
- Modularity
- High performance
- Async request handling

CONS

- Tailored to HEP?
- Often bolted on to storage system
- Often inefficient on small files
- Complexity



SRM global count

- DPM (Disk Pool Manager) - 144
- dCache - 59
- StoRM - 21
- CASTOR - 15
- BeStMan - 1
- Other or unknown - 12

Total in WLCG: 252



SRB

PROS

- Widely used in non-HEP
- Provides everything: a “data grid”
- Drag-and-drop interface (albeit limited)

CONS

- Only one implementation
- Not built for large files or high performance
- Closed source
- Proprietary protocols
- Different versions don't interoperate!



SRB National Count

- NGS has four SRBs at core sites
 - STFC, Manchester, Oxford, Leeds
- Other data storage/management options for NGS users
 - Disk!
 - Oracle databases
 - OGSA-DAI
- NGS (ETF) considering distributed filesystem between sites (DEISA-style)



iRODS

PROS

- Open source
- Partial re-engineering of SRB
- Flexible rule-based engine
- Complete “data grid”

CONS

- Only one implementation
- Proprietary protocol
- Doesn't fix all SRB's problems
- Needs recompiling when adding μ services
- Not production ready



Support

- GridPP supports DPM and dCache (and CASTOR^[*]) (and StoRM^[**])
 - 67 distinct subscribers (or thereabouts)
 - Based in 7 different countries
 - Core of ~12 regularly attending meetings
- Supporting middleware
 - Remit expanded to cover hardware
 - Including Data Management in GridPP3 (sort of)



(Potentially) Common Stuff

- Helpdesk (see all other talks 😊)
- Information services (SRM only)
- Fabric
 - Storage Hardware
 - Filesystems
 - Networks
- Distributed filesystems
 - Lustre, AFS, GPFS, ...

Procurement
Support
Tuning/Optimisation
Tech, and New Tech



Potential Requirements

- Small file handling
- Security
- Capabilities
 - Access latency
 - Retention policy
 - Storage type (eg custodial)
- Access interface
 - Web, Drag and Drop
 - POSIX
- Interoperation?
- Virtualisation?
- Scalable
 - Number of users
 - Freq, requests/sec
 - Number of files
 - Size of files
 - Total volume
- (To name but a few)
 - (Have more!)



Suggestions

- For now...
 - Share experiences when it makes sense (fabric)
 - Procurement, support, tuning, tech watch
 - Middleware: No one size fits all (users)
 - (Probably) raison d'être for all middleware
- Ponder requirements
 - Are they valid? Are they met?