WP5

Mass Storage Management

J Jensen
j.jensen@rl.ac.uk

DataGrid is a project funded by the European Commission
under contract IST-2000-25182

Outline

- Objectives
- Achievements
- Lessons learned
- Future & Exploitation
- Questions

Objectives

- Develop uniform interfaces to mass storage
  - Control interface
  - Data Transfer interfaces
  - Information interfaces
- Develop back-end support for mass storage systems
  - Provide “missing” features, e.g. directory support
  - Provide Grid access control
- Integrate with EDG Replica Management services
  - “Normally” users access SE via RM
- Publish information

Objectives

- Control interface
  - Original objective was “develop uniform interface to mass storage”
  - GSI authentication required
  - Interface changed to be a web service for compatibility with other WPs halfway through the project.
  - SRM version 1 was adopted as an alternative API for compatibility with other projects and LCG
- Data Transfer interface
  - Globus GridFTP required
  - Must support both encrypted and unencrypted transfers
- Information interface
  - Publish to MDS and, later, R-GMA

Achievements – Storage Element

- EDG Storage Element
- Flexible architecture
  - Cope with changing requirements
  - Pluggable features such as access control
  - Easy to add support for new storage systems
  - Easy to add new interfaces
- Security
  - Encrypted Globus GSI-authenticated control interface
  - File level access control (not in EDG 2.1 though)
- Currently supports CASTOR, HPSS, ADS, as well as disk

Achievements – site specific

- CASTOR SRM
  - Provided an SRM interface to CASTOR at CERN
  - Interoperability demonstrated with FermiLab
  - SRMCopy implemented
- CASTOR GridFTP
  - Provided a GridFTP interface to CASTOR's cache
  - Based on the Globus wu-ftpd GridFTP server
  - Files must be staged in before access
  - Transfer rates up to 30 MB/s (with specially tuned TCP settings)
- SARA
  - Porting SE to Irix, developing cache management tools

Achievements – collaborations

- Contributions to international standards and fora
  - SRM
    - Collaboration between Fermilab, Jefferson Lab, Lawrence Berkeley, RAL, CERN
    - Contributed to the design of the SRM version 2 protocol
  - GLUE
    - Contributed to the design of GLUE storage schema
  - GGF
    - Tracked developments in appropriate working groups
  - Dissemination
    - Talks at conferences and in working groups, publications,…
- EDG
  - Participated in ITeam, ATF, SCG, QAG,…

Lessons learned

- Choice of architecture was definitely right
  - Architecture has successfully coped with changing requirements
- Look for opportunities for component reuse
  - Used web services deployment and security components provided by WP2
  - Deployed and developed further information producers supplied by WP3
  - Almost all Data Transfer components developed externally
- Scalability
  - Use on application testbed or in production found scalability problems not seen on development and integration testbeds

Lessons learned

- Inter-WP integration requires a lot of effort !
  - At times, nearly 100% of WP5 devoted to ITeam work and site installation support
  - Storage interface machines are heterogeneous
    - More installation support was required
  - For example, effort required to support DICOM servers was significantly underestimated
    - Requires significant effort from WPs 2, 3, 5, 10 – plus of course SCG, ATF, and, eventually, ITeam
- Prototype implementations live longer than expected
  - SE's metadata system was implemented as prototype
  - Needs modifications to scale – first step (of two) almost complete
  - Second step : need to choose appropriate database

Exploitation

- Used yesterday in middleware demo to access mass storage
- Used successfully on all EDG testbeds by all EDG applications WPs at all participating sites
- "Atlas Data Challenge 1.5"
  - SE is currently used by Atlas to transfer data between ADS at RAL and CASTOR at CERN
  - About 1500 files; 2 TB in total
  - Files are copied by EDG RM and registered in an RC at RAL
  - This work is being done by Atlas outside the EDG testbeds
- Storage Element information producers
  - Information producers used independently, e.g. by LCG and DZero in the "classic" SE

Future and exploitation

- Storage Element SRM
  - SE will provide generic SRM 1 interface
  - This work is almost finished
  - Work will be carried on by RAL; later in GridPP 2
  - To be tested by LCG
  - Will investigate whether to build SRM version 2
    - Depends on uptake of protocol in international community
    - Current SRM implementation is built with also SRM 2 in mind
    - Some additional features required
- Storage Element – further mass storage systems
  - Scope for implementing support for AMS, DICOM?
  - Support for UK Tier-2 sites to be developed by GridPP2

Future and exploitation

- Storage Element and VOMS
  - Integrate VOMS support into SE – SE already works with VOMS proxies
  - Will enable more scalable access control
  - Fairly easy task – accomplished again by reusing components
  - May need to VOMS-enable GridFTP server – integrate LCAS and LCMAPS
- Integration with GFAL
  - LCG's Grid File Access Library – POSIX style interface
  - Planned integration using SRM 1 interface
- Automatic Grid mirroring
  - Looking into automatic mirroring data in UK NeSC – between Edinburgh and Glasgow

Questions?

Achievements – Storage Element

Storage Element Component overview

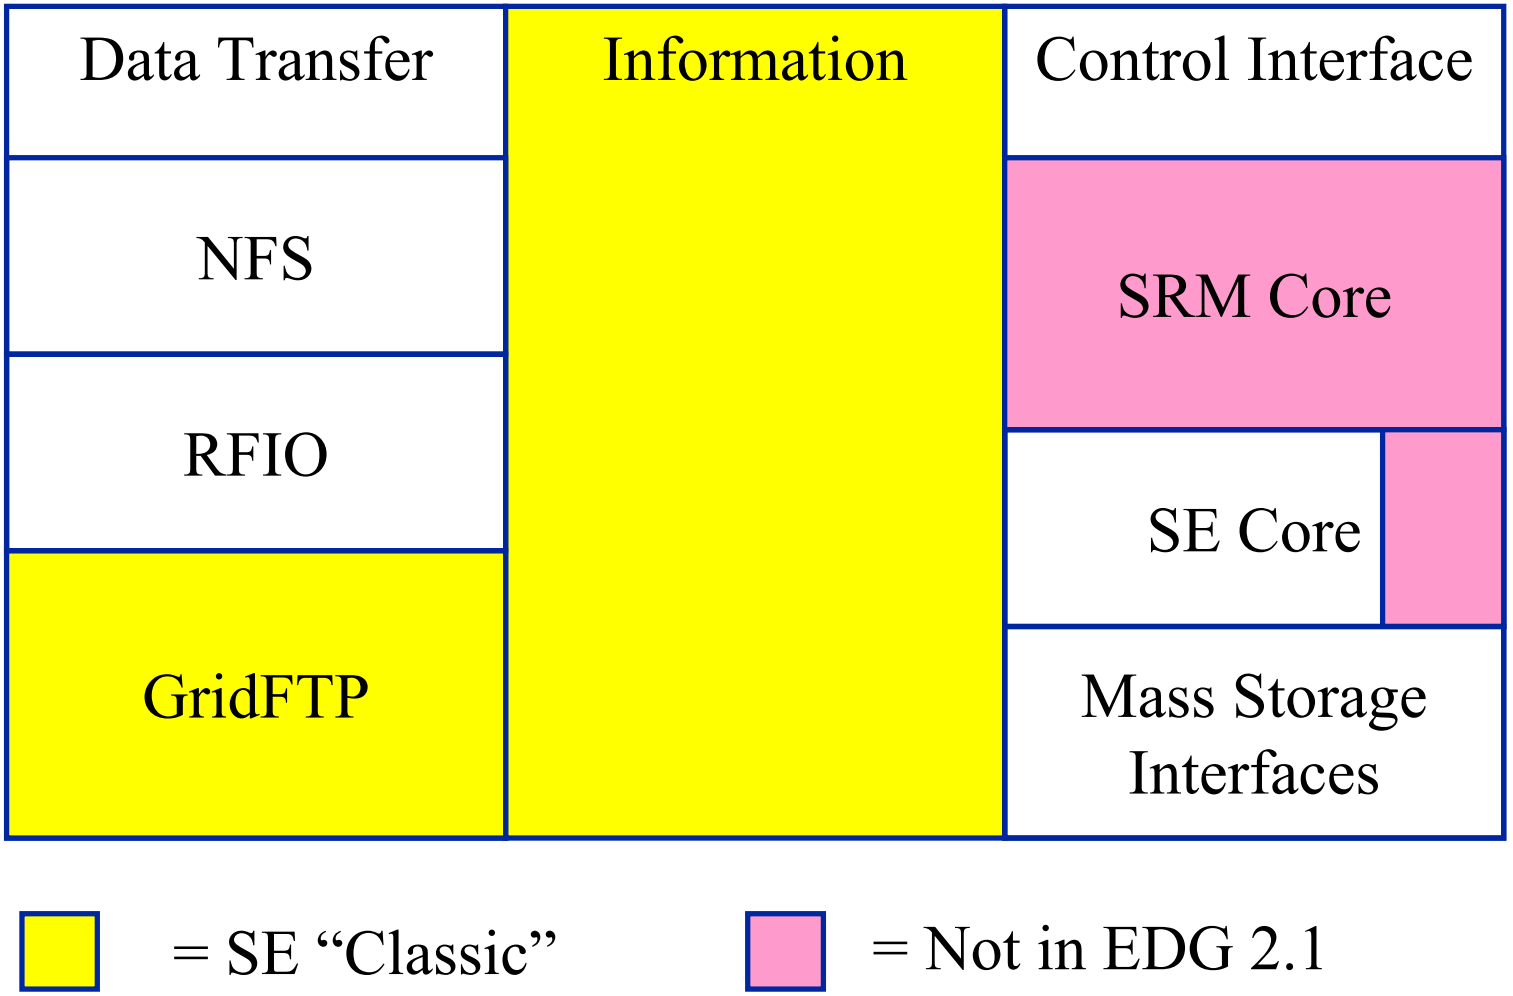

GFAL, SRM, and Storage Element

- LCG decided to use GFAL – the "Grid File Access Library"
- It was decided to interface to EDG SE using SRM 1 interface
- SRM 1 can also be used for interoperability with DoE Labs
- We are integrating the EDG SRM layer with the EDG SE
- Some complications → not in 2.1
- We are committed to completing the task