LCG Launch: Summary of Grid Technology (WG3)

Norman McCubbin (RAL)

Laura Perini (Milano)

(not Technology experts!)

+ much input from session speakers and others.

Grid Technology: Summary

- Agenda of 13 March session:
  - Introduction and Overview: Ian Foster
  - Data Management: Peter Kunszt
  - Security: Dave Kelsey
  - Scheduling (US): Miron Livny
  - Scheduling (EU): Francesco Prelz
  - Information Services: Schopf/Magowan
  - Discussion: led by Fabrizio Gagliardi

Grid Technology: Summary

- What is "Grid Technology"?
  - It's the bit between the user's or experiment's application and the (Grid-enabled) computer system ("fabric") at a particular institute or laboratory;
  - For users, it's the bit we don't want to have to worry about, provided it's there!
  - Note the analogy with electric-power grid: you plug your device into the socket in the wall; you do not care, or want to care, about the power stations, distribution system, HV cables,……

The LHC Computing Grid (LCG) Project Structure

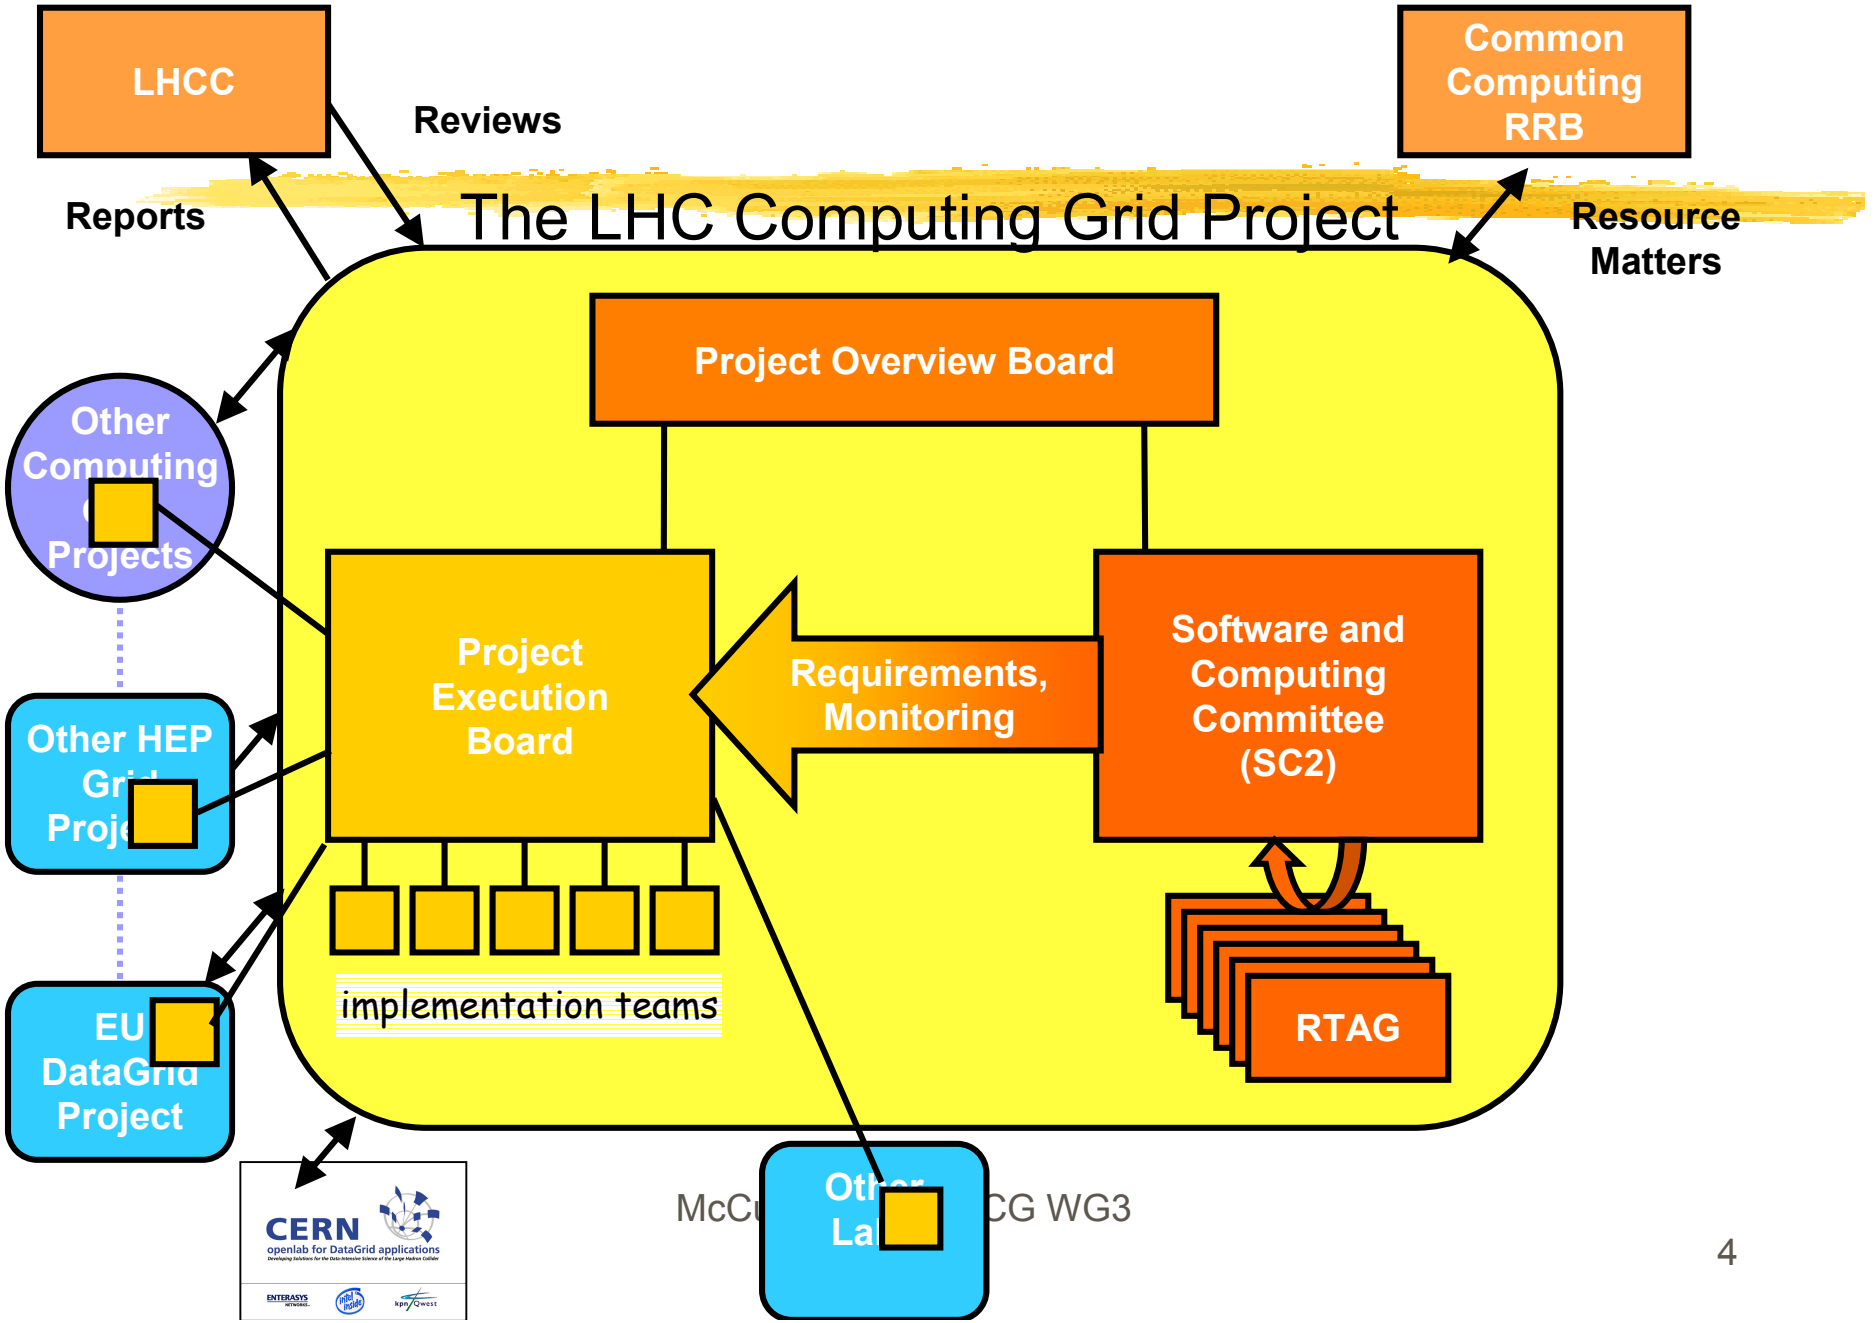

Grid Architecture (Foster)

"Coordinating multiple resources": ubiquitous infrastructure services, app-specific distributed services

"Sharing single resources": negotiating access, controlling use

"Talking to things": communication (Internet protocols) & security

"Controlling things locally": Access to, & control of, resources

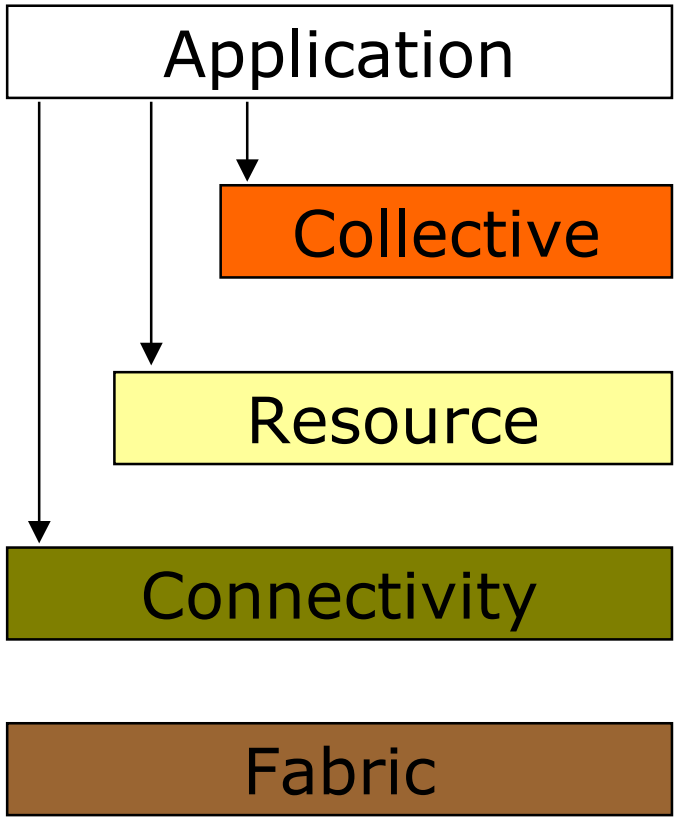

For more info: www.globus.org/research/papers/anatomy.pdf

The Grid Problem (Foster)

Resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations (VO)

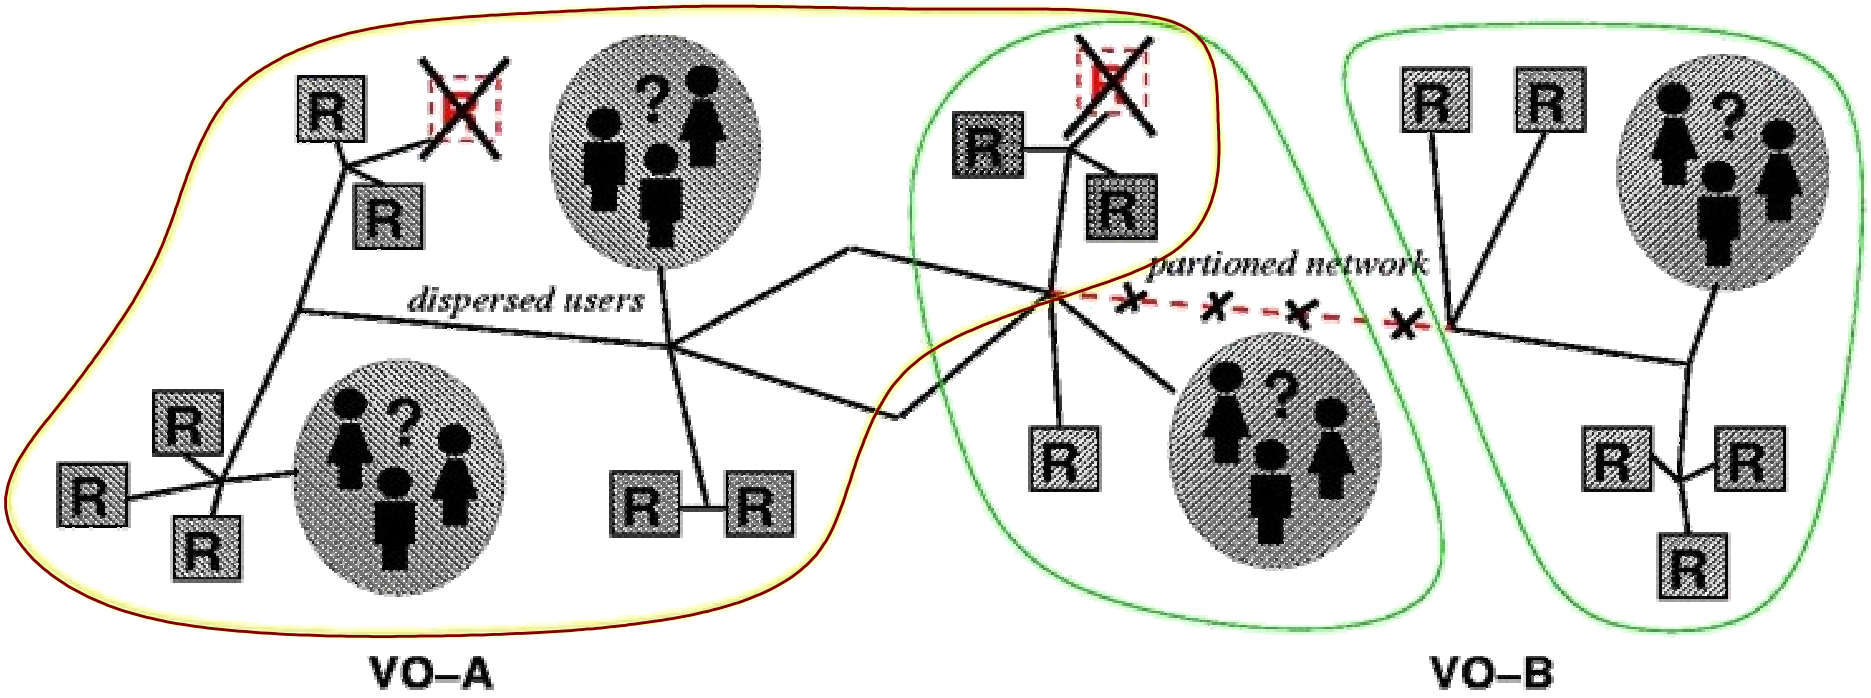

Aspects of the Problem (Foster)

1) Need for interoperability when different groups want to share resources
   - Diverse components, policies, mechanisms
   - E.g., standard notions of identity, means of communication, resource descriptions
2) Need for shared infrastructure services to avoid repeated development, installation
   - E.g., one port/service/protocol for remote access to computing, not one per tool/appln
   - E.g., Certificate Authorities: expensive to run

⌘ A common need for protocols & services

Web Services (Foster)

- "Web services" provide
  - A standard interface definition language (WSDL)
  - Standard RPC protocol (SOAP) [but not required]
  - Emerging higher-level services (e.g., workflow)
- Nothing to do with the Web (!)
- Useful framework/toolset for Grid applications?
  - See proposed Open Grid Services Architecture (OGSA)
- Represent a natural evolution of current technology
  - No need to change any existing plans
  - Introduce in phased fashion when available
  - Maintain focus on hard issues: how to structure services, build applications, operate Grids

For more info: www.globus.org/research/papers/physiology.pdf

What do users expect…?

Even if users/applications "don't want to know" the gory details, we expect:

- Data Management **across the Grid**;
- Efficient "scheduling";
- Access to information about what's going on, either on the Grid itself, or on the component computer systems (fabric layer).
- …and we know that someone will have to worry about SECURITY, but we don't want to!

Vision of Grid Data Management (Kunszt)

Ubiquitous Data Access ("AFS" on the Grid)

- Global Namespace
- Transparent security control and enforcement
- Access from anytime anywhere, physical data location irrelevant
- Automatic Data Replication and Validation

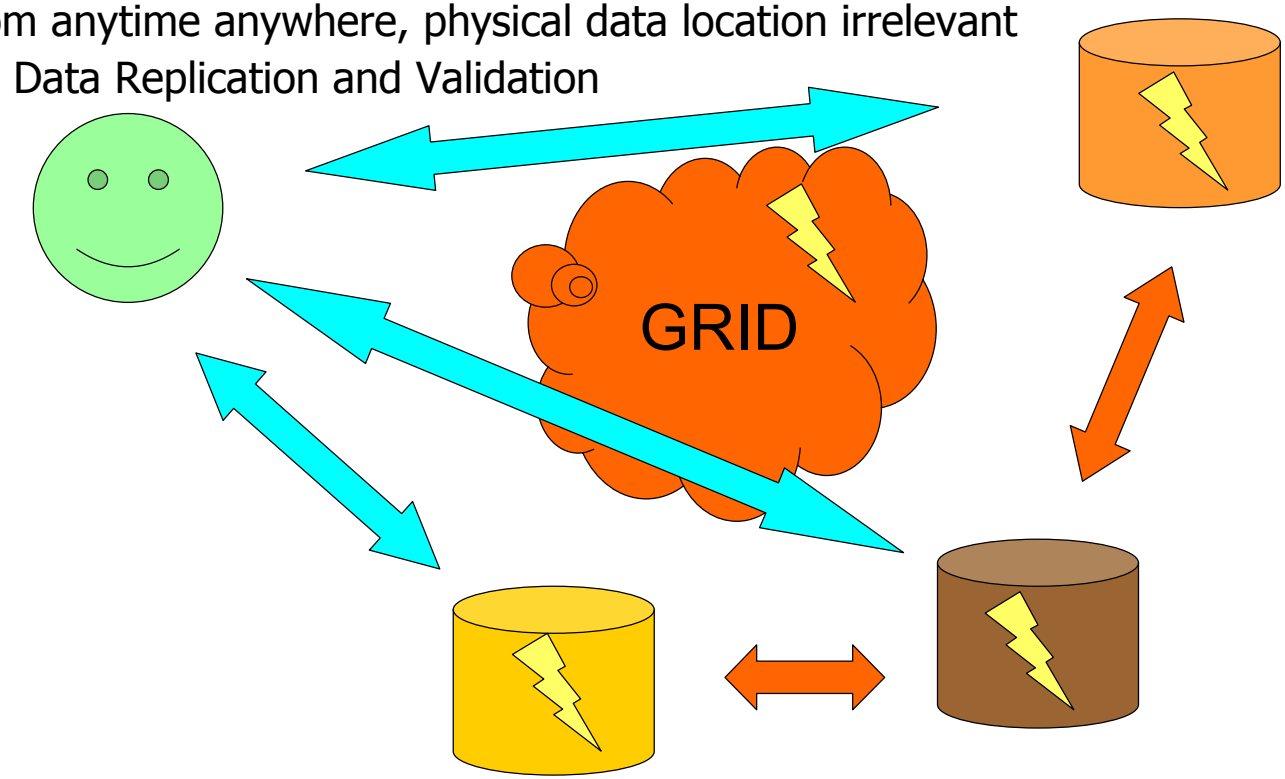

Data Management: EDG WP2 (Kunszt)

- **GDMP** – with PPDG
  - In production with CMS for Objectivity replication
  - Subscription-based replication
  - Scalable architecture
- **Replica Catalog** with Globus
- **Replica Manager** and Optimiser
  - Take Globus RM as core
  - Additional modules for pre- postprocessing of data
- **Replica Selection** in the WP2 Optimisation task
  - Simulator to test replica selection
- **Spitfire**
  - Unified front-end to databases
  - Suitable for Grid and Application Metadata

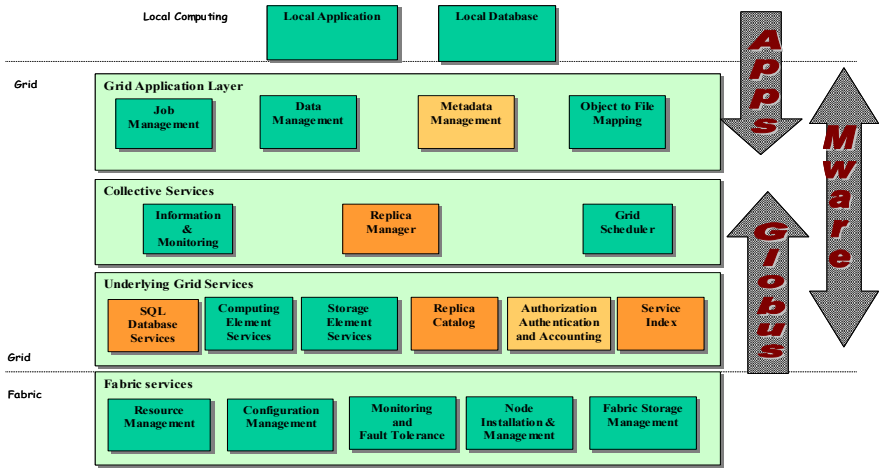

Data Management: PPDG / Griphyn (Kunszt)

- Globus, Condor, SRB
- GDMP – with EDG
- Magda
  - To be used in ATLAS data challenges
  - Metadata catalog
- JASMine JLAB Asynchronous Storage Manager
  - Storage Management and Resource
  - Replica catalog based on MySQL, as Web Service
  - Replication service
  - File Server
- Griphyn Virtual Data System

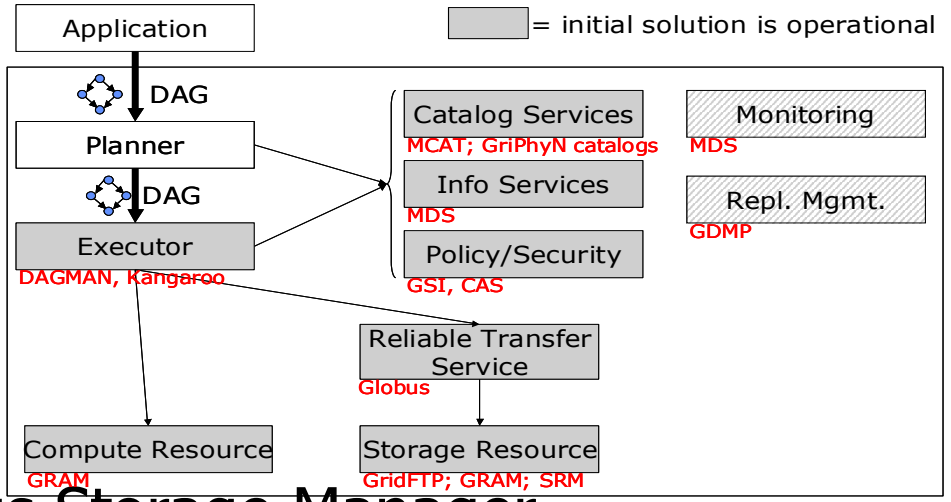

Issues / Dangers (Kunszt)

- Commonalities – solving the same problems again and again ; potential for duplication of effort
  - Think in Virtual Organisations (VO)
  - RTAGs, like Common Persistency Framework
- Security – i can see what you can't see
  - EDG Security Group – see Dave Kelsey's talk
  - SciDAC
  - Building Trust relationships
- **Standardisation** – bringing it all together and **agree, agree, agree**
  - OGSA ("Web Services")
  - GGF
- Consensus – too many cooks spoil the broth
  - Making decisions in time
  - **Keeping agreements, sticking to standards**
  - Avoid Micromanagement

“Scheduling”: Driving Concepts (Livny)

- Virtual Data – Service Requests are in the form
  “ place y = F(x) at location L”.
  I.e. you want to do something (F) on data (x) to obtain a result (y); the result is wanted at a particular location (L), e.g. your lap-top;
- A common reference Data Grid Architecture.
- Integrated research efforts in the areas of planning algorithms and scheduling policies.
- A framework of networked services connected by reliable, recoverable and bi-directional interfaces.
- Uniform view of processing and data placement activities.
- Job flow management based on **Directed Acyclic Graphs (DAGs)** of jobs.
- **Evaluation of Grid technology via end-to-end implementations of “real-life” services.**
- Interoperability with “external” technologies and deployed infrastructure.

“Scheduling”: Services (soon) Available (Livny)

- The Globus Tool Kit – **Inter**-domain information security and job submission services.
- The Condor system – **Intra**-domain information, security, job management and resource allocation services.
- Condor-G – Job management services for Globus jobs
- RLS – Logical to physical mapping of file names
- DRM – Data staging services
- VDS – Virtual Data Language, Derivation and Transformation Catalogs
- GridFTP and RFT – file/data movers
- DAGMan – Job flow services

“Scheduling”: EDG (Prelz)

- **An incremental approach** EDG WP1 delivered (and is currently supporting) the following functionality for the first project release:
- Ability to **submit a job** (described via the Condor ClassAD-based Job Description Language, or JDL) to the DataGrid testbed from any user machine.
  - Lightweight, python-based client, with a dependency on Globus GSSAPI, and thus on OpenSSL.
- The WP1 client allows to monitor and control (terminate) the job, *and* to transfer a "small" amount of data to and from the client machine and the executing machine.
- WP1's **Resource Broker** chooses an appropriate computing resource for the job, based on the constraints specified in the JDL.
  - where the submitting user has proper authorization
  - that matches the characteristics specified in the job ClassAD (Architecture, computing power, application environment, etc.)
  - and where the specified input data (and possibly the chosen output SE) are determined to be "close enough" by the appropriate resource administrators.
- Throughout this process, WP1's Logging and Bookkeeping services maintain a "state machine" view of each job.

“Scheduling”: EDG (Prelz)

- **On-going developments for EDG Release 2**
- Integration of WP2 "query optimisation" (based on network information and driving data replication).
- Support (and API) for the optimized scheduling of partitionable and checkpointable (embarrassingly parallel) jobs.
- Submission of multiple (and possibly dependent) jobs - DAGMan.
- Provision of APIs for the applications.
- Prototype GRID accounting system (based on an economic model).
- Support for "interactive" jobs (as defined/required by the applications).
- Support for single-cluster MPI jobs.
- Development of GUI components.

Information Services: Metacomputing Directory Service (MDS) (Schopf)

- Globus Information Service (GIS)
- Used by iVDGL, GriPhyN, PPDG, EDG, NMI, Grads, etc.
- Requirements and characteristics
  - Uniform, flexible access to information
  - Scalable, efficient access to dynamic data
  - Access to multiple information sources
  - Decentralized maintenance
  - Secure information provision
- Main contact John McGee (mcgee@isi.edu)

Information Services: Grid Monitoring Architecture (GMA) (Magowan)

- We use it not only for monitoring but also as the basis of an information system

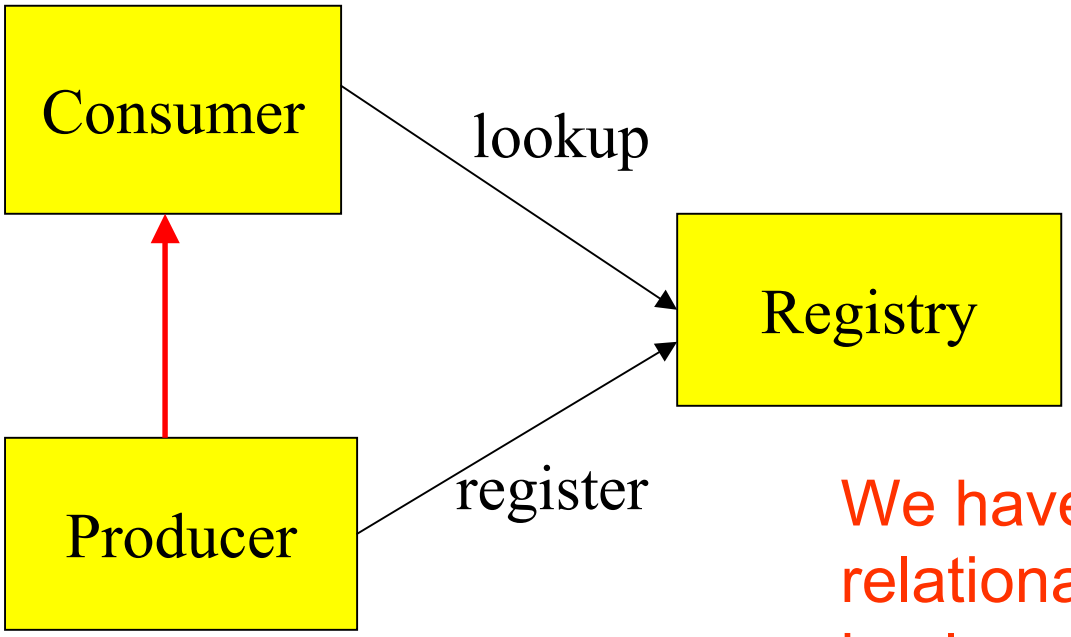

We have chosen a relational implementation

Information Services: Co-ordination Efforts

- How to co-ordinate?
  - Within the US: Joint Monitoring Project PPDG/GriPhyN/iVDGL;
  - International I: HICB/JTB (see later) "GLUE" effort;
  - International II: Discovery And Monitoring Event Description (DAMED) Working Group.
- (Who co-ordinates the co-ordinators??)

Security: Overview (Kelsey)

- Security requirements
- AAA Architecture (*Authentication, Authorisation, Accounting*)
- Technology and Grid projects
  - Globus
  - DataGrid
  - PPDG
  - DataTAG/iVGDL/HICB
  - SecureGRID
- Security Issues
  - Authentication
  - Authorisation
  - Grid Deployment

Security Requirements (Kelsey)

The usual tension: *functionality vs. security*

- But with some special features
  - Scale of users and resources

- **Site Security Officer**
  - Protect the site from hostile attack
- **Resource/Site System Manager**
  - Complete control of the local resources
- **Virtual Organisation**
  - Allocate resources to members, groups, roles
- **User**
  - Easy and transparent access to resources

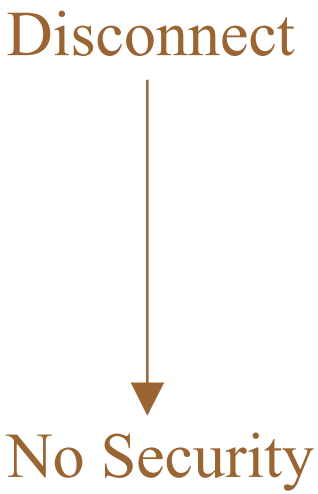

Security: Authorisation issues (Kelsey)

- In contrast to Authentication, the technology for Authorisation is much less mature.
- NB: Some users will belong to multiple VO's
  - Authorisation may need to be based on "joins"
- Global vs Local authorisation mechanisms
- We need more functionality
  - "Dynamic policy-based Access control"
  - Users with more than one allowed role
  - Move away from Unix uid based security? (and grid mapfile?)
  - Applicable to all Grid services (and callable from)
    - Maybe different levels for different services
  - need to negotiate policy – Global/VO/Local

Security: Co-ordination: DataTAG/iVDGL/HICB

- Transatlantic Testbed(s)
  - Interoperability essential for LCG applications!
- Cross project Authentication
  - US DOE SciGrid CA already "trusted" by EDG
  - US projects working on "trust" of EDG CA's
- Cross project Authorisation
  - DataTAG WP4 has resources to work in this area

Grid Technology: Issues (1)

- Co-ordination of Grid efforts: Users/experiments must see common world-wide **interface**; Grid Technology must deliver this. The various Grid projects recognise clearly the importance of this. (See for example previous transparencies)
- Note that more than one approach in Grid Technology is **healthy** at this stage. It takes time to learn what should be common!
- Note there is BOTH Global Grid Forum (GGF) for **all** Grid activity, not just HEP, and also **HICB/JTB** which addresses co-ordination for HEP and Nuclear Physics. HICB started about a year ago and has recently proposed "GLUE" (Grid Laboratory for a Universal Environment) initiative to select (develop) a set of Grid services which will inter-operate world-wide. GLUE is based on iVDGL (US) and DataTag (EU).

Grid Technology: Issues (2)

- If HICB did not exist, LCG would have to do this for LHC. LCG still has responsibility for defining the Grid requirements of the LHC experiments ("RTAG" about to start), but will obviously evaluate and hope to use GLUE.
- LCG should, nevertheless, keep a watchful eye on the various 'middleware' projects to minimise risk of divergences.
- LCG does not expect to devote any of its own resources to middleware development. Les Roberston: "LCG is about Grid **deployment**." As a "consumer" of Grid Technology, LCG should be concerned about the "supply": hence continued **financial support** for EDG, PPDG, GriPhyN… and their **successors is vital.**

Grid Technology: Issues (3)

- It is absolutely **essential** that the LHC experiments test, as realistically as possible, the Grid Technology as it becomes available. This has already started, but will increase, particularly in the context of Data Challenges. This is vital not just for the usual reason of incremental testing, but also to discover the **potential of the Grid for analysis**. (Rene Brun)
- This 'realistic' testing generates **"tension"** between **developing** new features of Grid Technology, and **supporting** versions used by the experiments.

Grid Technology: Issues (4)

- We can expect a certain level of support from the middleware providers (the Grid Projects), but LCG should **expect to devote some resources** in this area: both for middleware installation and a "help-desk" as an interface between middleware users and the developers. (Note importance also here for continued funding for the middleware developers.)
- Although finally users shouldn't have to care about the Grid Technology layer, during the development phase the experiments (through the LCG) should ensure that the appropriate interfaces between the experiments' software and the Grid Technology layer(s) are developed. This will require effort from **both middleware and LHC software developers (architects).**