

Review of the LCG

Ian Fisk
DOSAR Workshop
March 30, 2006

History



The LHC Computing Grid (LCG) Project was formed in January of 2002

➡ It consisted of 4 areas

- Grid Development
- Grid Deployment
- Application Area
- Fabrics

| What we normally think of as LCG the Grid

➡ The project was well funded and well staffed

- The first two areas are what is most associated with LCG externally, but the funding had large contributions to the Fabric (Facilities and upgrades at CERN) and the Applications areas

- LCG-AA was responsible for core common projects used by the LCG experiments
- POOL the persistency mechanism, shared library SEAL, etc.

➡ Grid deployment operates the LCG computing grid infrastructure.

➡ Grid development is responsible for the service development

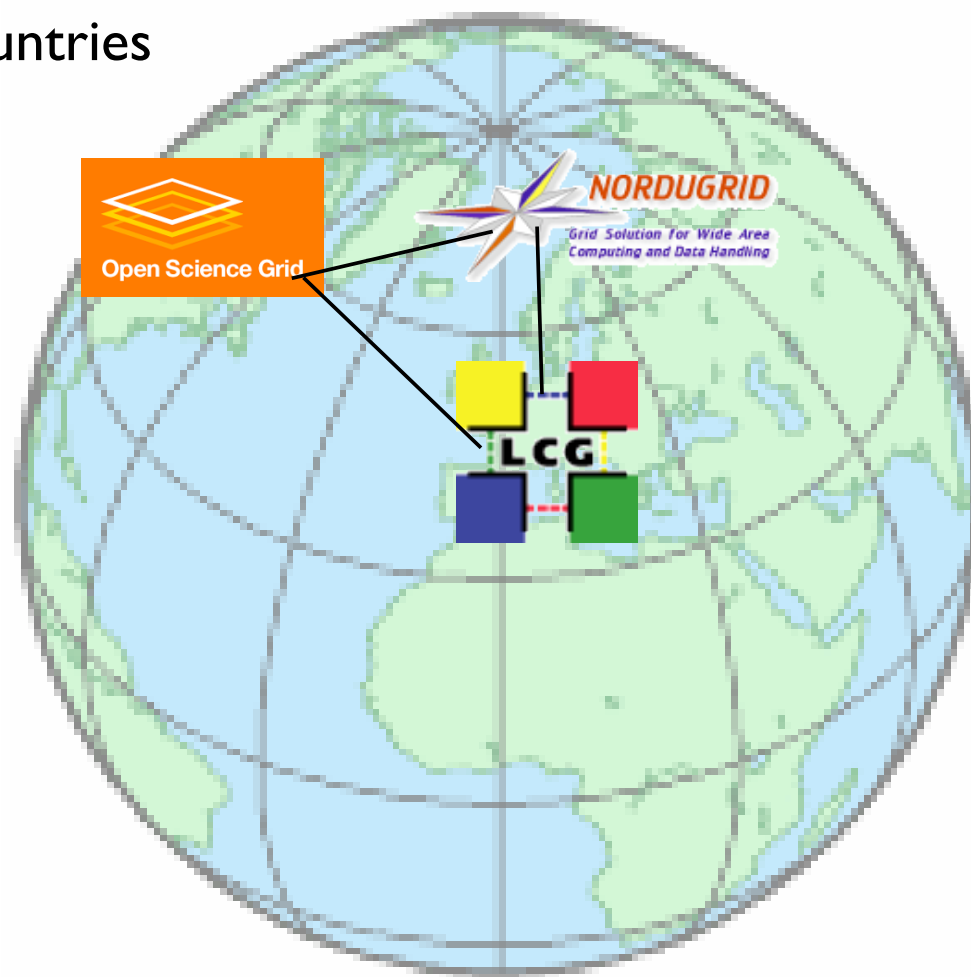
The Worldwide LHC Computing Grid



In 2005 the WLCG was formed. It is the Federation of 3 grid infrastructures

- ➡ LCG in most of Europe
- ➡ NorduGrid in the Nordic Countries
- ➡ Open Science Grid in the Americas

Today I'll focus on LCG the grid infrastructure project, which is supported by EGEE

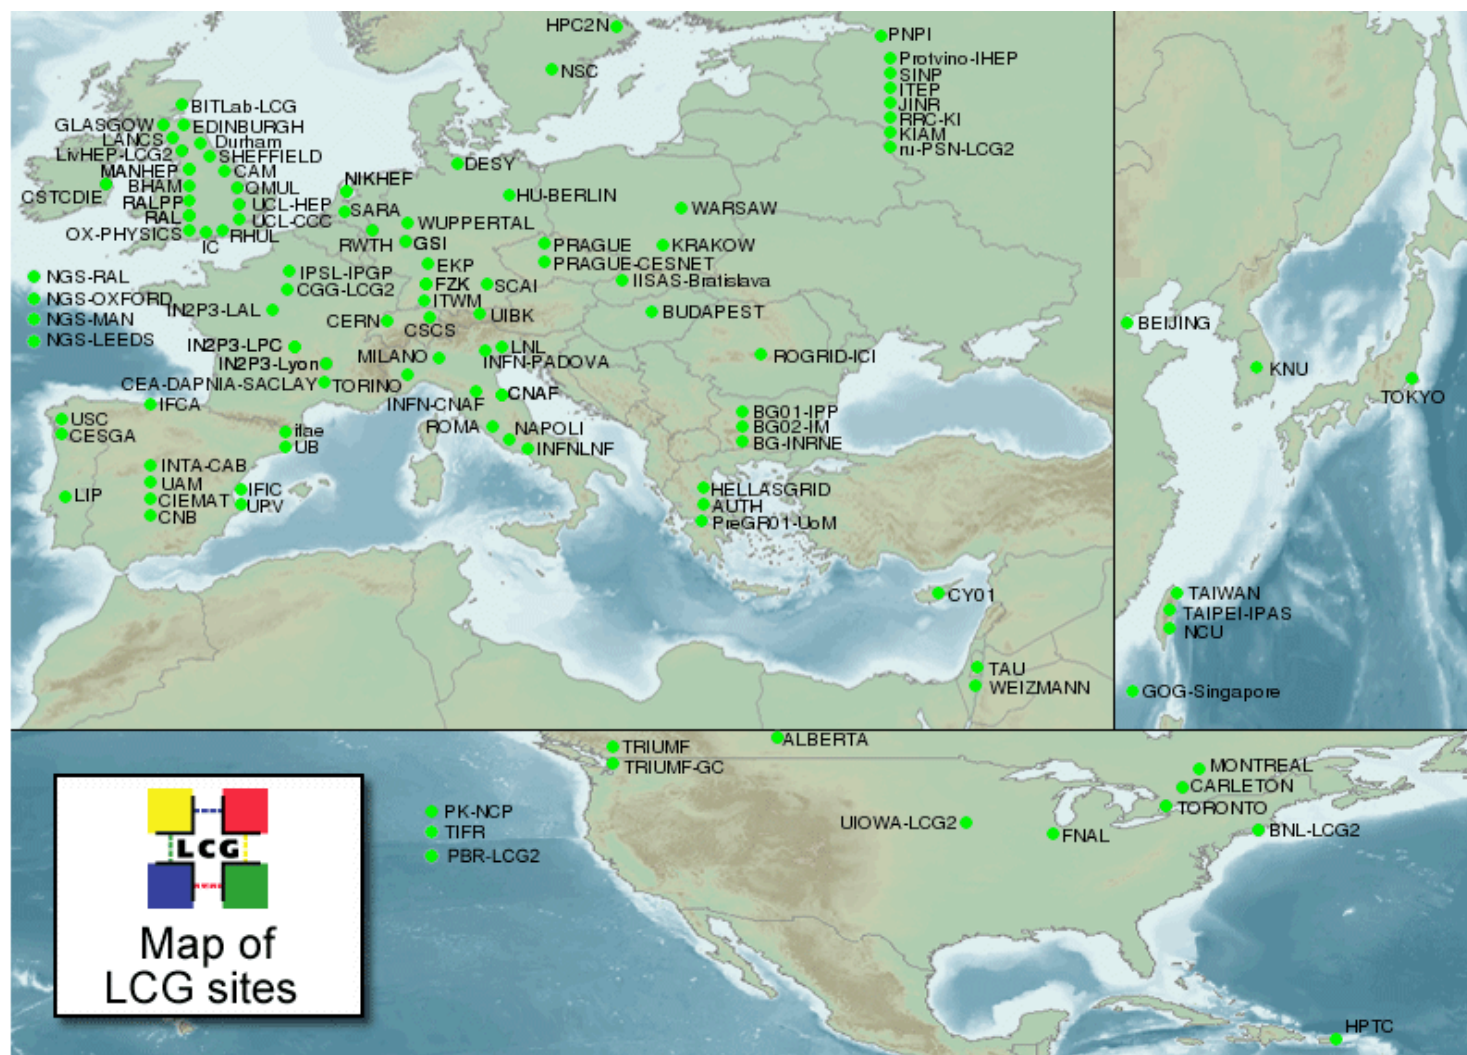


Scale and Distribution




More than 100 sites in 31 countries

10k CPUs, Petabytes of total disk,



Supported Virtual Organizations



There are more than 15 Virtual organizations supported by the LCG infrastructure

- ➡ Though most of the activity is related to LHC preparations
 - ATLAS, CMS, ALICE, and LHC-B have all made significant use
- ➡ There are a number of high energy physics VOs
 - Babar, D0, CDF are all registered
- ➡ There are some VOs devoted to software development
 - GEANT4
- ➡ And some organizations from other sciences
 - Biology and Earth Science

Underlying Components



The principle components of the LCG middleware are

- ➔ The Globus Toolkit (GT2) developed by the Globus Project
- ➔ The Condor system developed at the University of Wisconsin, Madison
- ➔ The Globus and Condor components and some other tools from US projects are integrated and packaged as the Virtual Data Toolkit by the VDT project at the University of Wisconsin, Madison. VDT provides support for this package to LCG/EGEE.
- ➔ Tools developed by the DataGrid Project (EDG). The EU-funded DataGrid project ended in 2004, but the institutes that had developed the tools needed for the LCG/EGEE grid continue to support them until they are replaced by improved software.

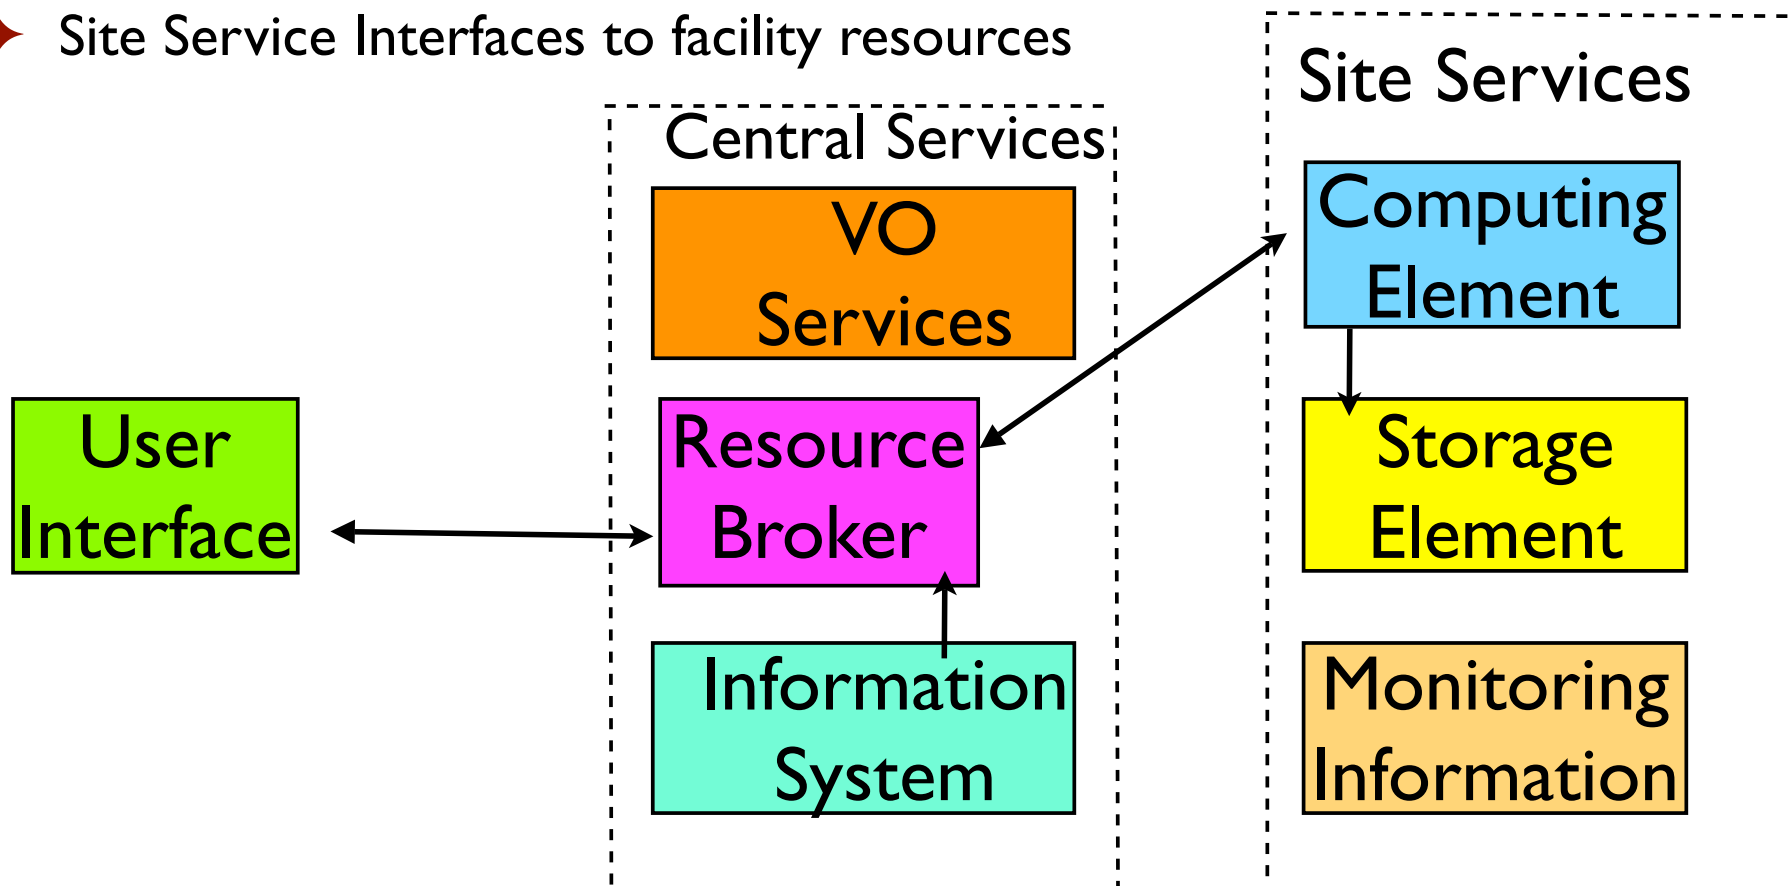
New middleware components developed as part of the gLite toolkit by the EGEE project. The first release of gLite will provide improved tools for workload scheduling, grid catalog, and a monitoring infrastructure.

Basic Connections



Services are roughly divided into three area

- ➔ The User Interface are installed for user submission
- ➔ Central services for brokering, VO organization and information collections
- ➔ Site Service Interfaces to facility resources





User Interface

The User Interface (UI) is the client submission tools for the Grid

- ➡ Services for authentication
 - Grid-proxy-init and voms-proxy-init
- ➡ Job submission and tools to return results
 - Normal globus submission, but European data grid (edg) job submission tools to submit to the Resource Broker and to retrieve job output
 - Access to lcg client tools for data registration and retrieval
- ➡ Query tools against the central logging and bookkeeping service

Central Services



VOs are registered and organized with a product called VOMS

Central Services

VO
Services

Resource
Broker

Information
System

- ➔ allows formation of roles and groups

One of the big additions to grid technology provided by the LCG is the resource broker (RB)

- ➔ Jobs are submitted to the RB, which handles the input and output sandbox
 - RB can make decisions for submission based on site name, VO constraints like installed software, and dynamic information like system load

RB needs a lot of configuration information to make meaningful scheduling decisions

- ➔ Information is collected in a scalable implementation called the BDII
 - Collect from site GRISs or other BDIIs

Site Services



Site Services

Computing
Element

Storage
Element

Monitoring
Information


These are interfaces to the facility resources

- ➔ Computing Element is the interface to the batch system
 - LCG has made improvements in the scalability of the job managers
- ➔ Storage Element is the interface to disk, tape, or hierarchical mass storage (CASTOR, dCache DPM)
 - Storage Resource Manager (SRM) has been chosen as the interface to storage
- ➔ Monitoring, Information, and Accounting are collected with a variety of services
 - Information system is done with a local BDII
 - The accounting system is an area of significant development for LCG
 - Many LCG sites support multiple experiments


LCG-2 Grid Use



Jobs submitted by VO (longer than 10 minutes)



Number of individuals submitting jobs



New Development



EGEE is in the process of developing and testing the gLite 3.0 software release


The EGEE software stack has been in development for over a year

- ➡ The current version, which is being rolled out is a combination of the functional versions of the previous release and the new components of gLite.

Improvements in the architecture and scale of a number of existing components

- ➡ Current RB is quite functional but very slow for submission
 - gLite RB enables bulk submission, which should improve performance
- ➡ Catalog used for registering output files when they enter the storage element has been reworked for improved performance and scaling
- ➡ Additions to monitoring and informations systems
- ➡ Additions to the accounting services

New Development (Continued)



In addition to the improvements of the existing services there are new services being rolled out for gLite

- ➔ File Transfer Service transfers files between storage elements
 - designed to improve the reliability to move files between sites
 - tested with promising results during the rerun of the SC3 throughput phase
- ➔ The gLite Workload Management System is an implementation of Condor-C and implements some of the workflows currently handled in the RB
 - Should allow an improved scale

LCG and OSG interoperability



The LCG grid deployment activity and the Open Science Grid have been working on interoperability with reasonable success

- ➡ Building on VDT and common low level components
- ➡ Building on common efforts like the GLUE schema to describe the information about a site
- ➡ Common services
 - VOMS for managing virtual organizations
 - VDT components

We now have the ability to do cross grid submission

- ➡ Still requires some effort to keep the submission functional
- ➡ Work left to do, but the federated grid to support experiments is coming
 - Working on accounting, monitoring, auditing and security issues
 - A number of the technical issues for submission have been solved.