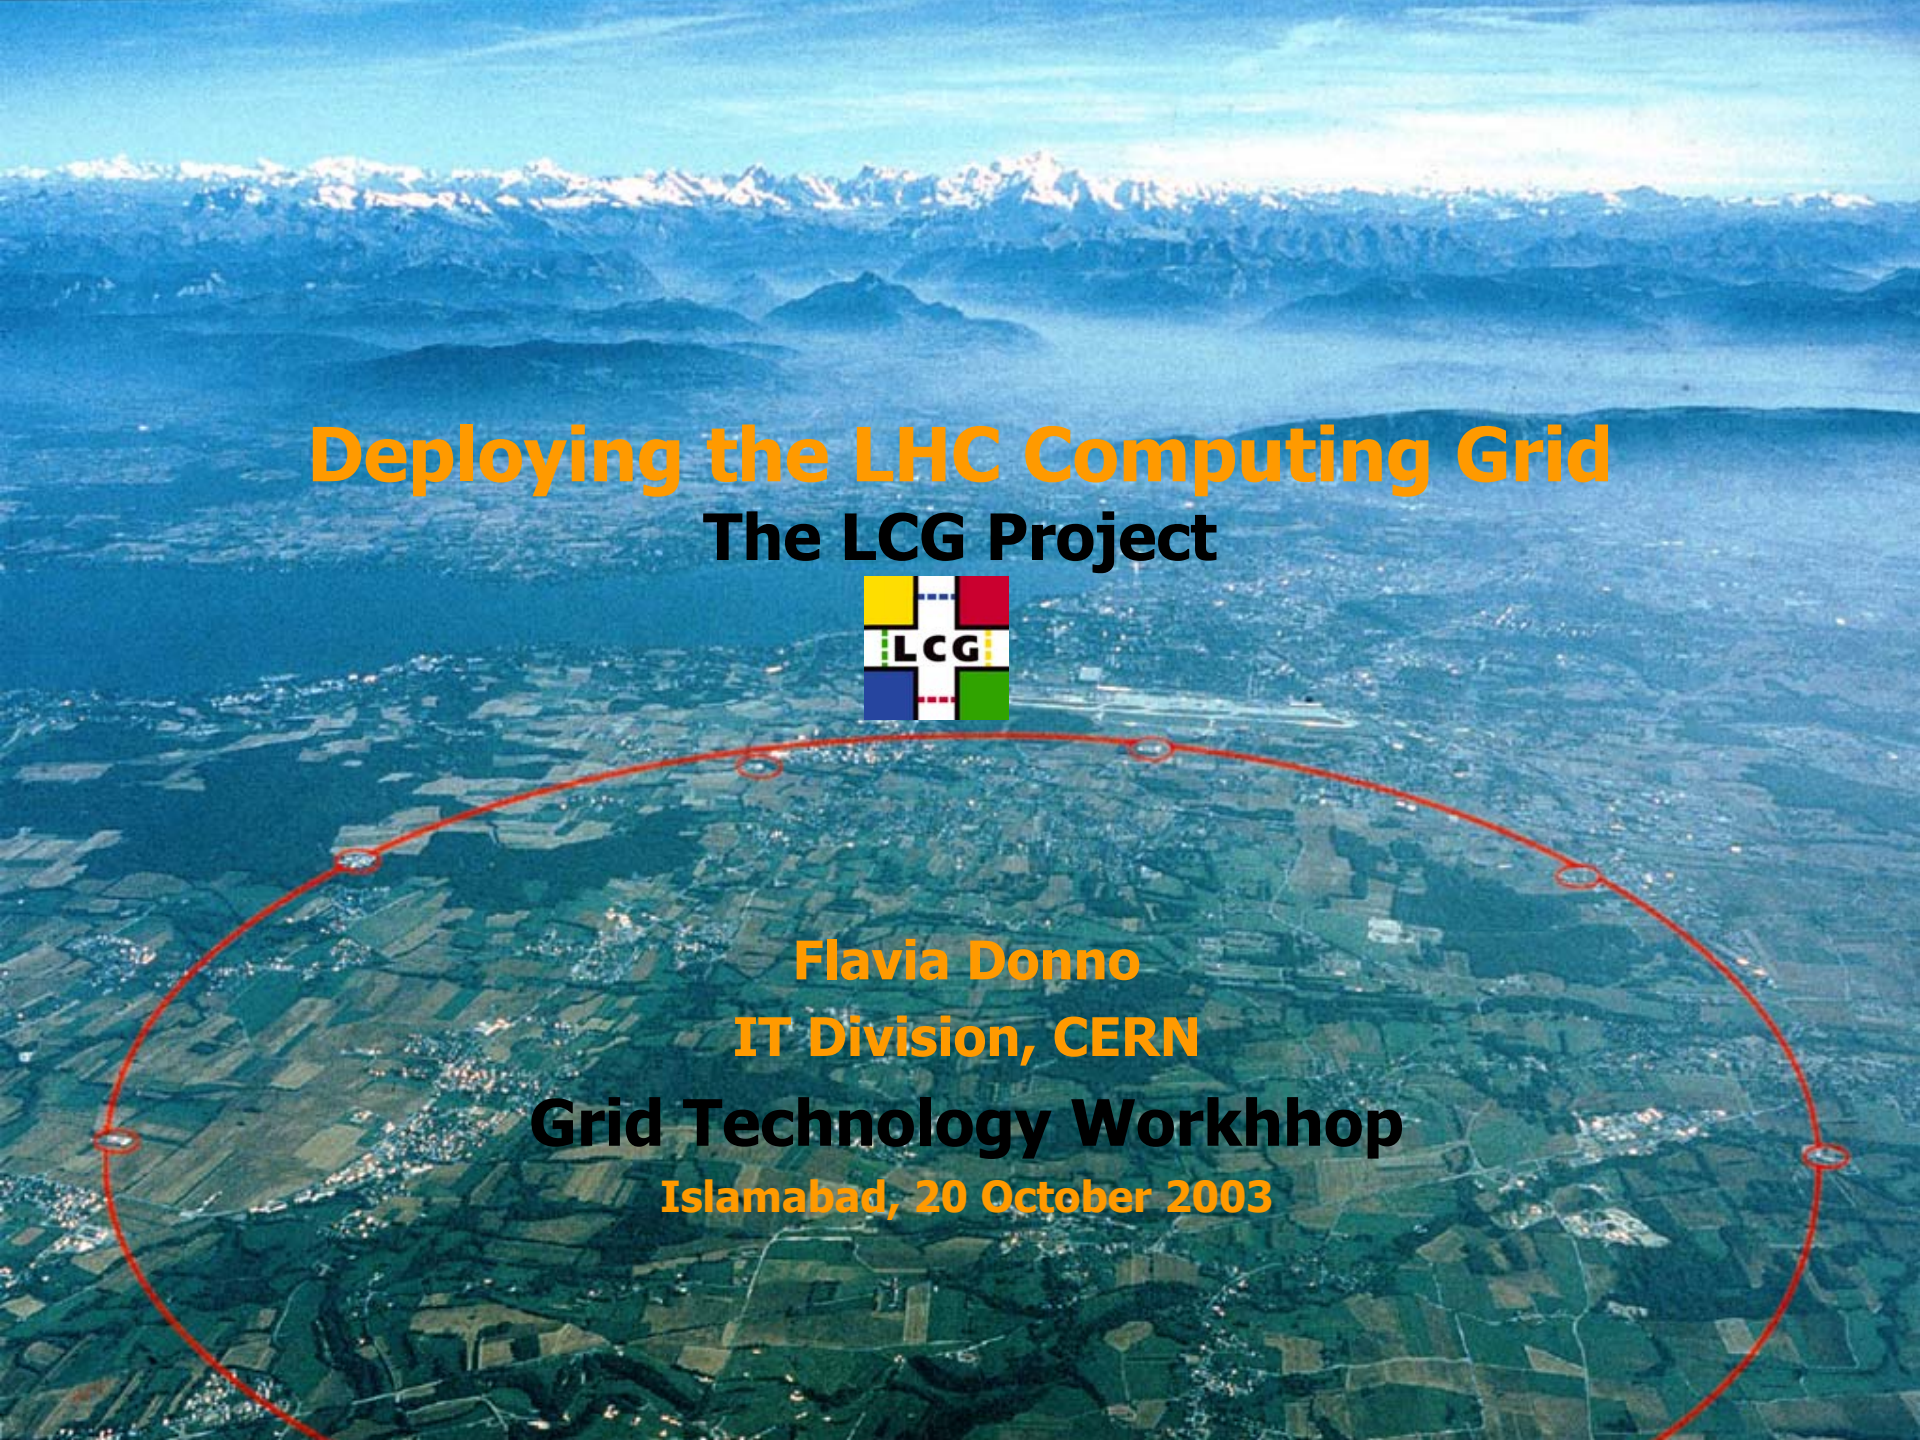

Deploying the LHC Computing Grid

The LCG Project

Flavia Donno

IT Division, CERN

Grid Technology Workhhop

Islamabad, 20 October 2003

The Large Hadron Collider Project 4 detectors

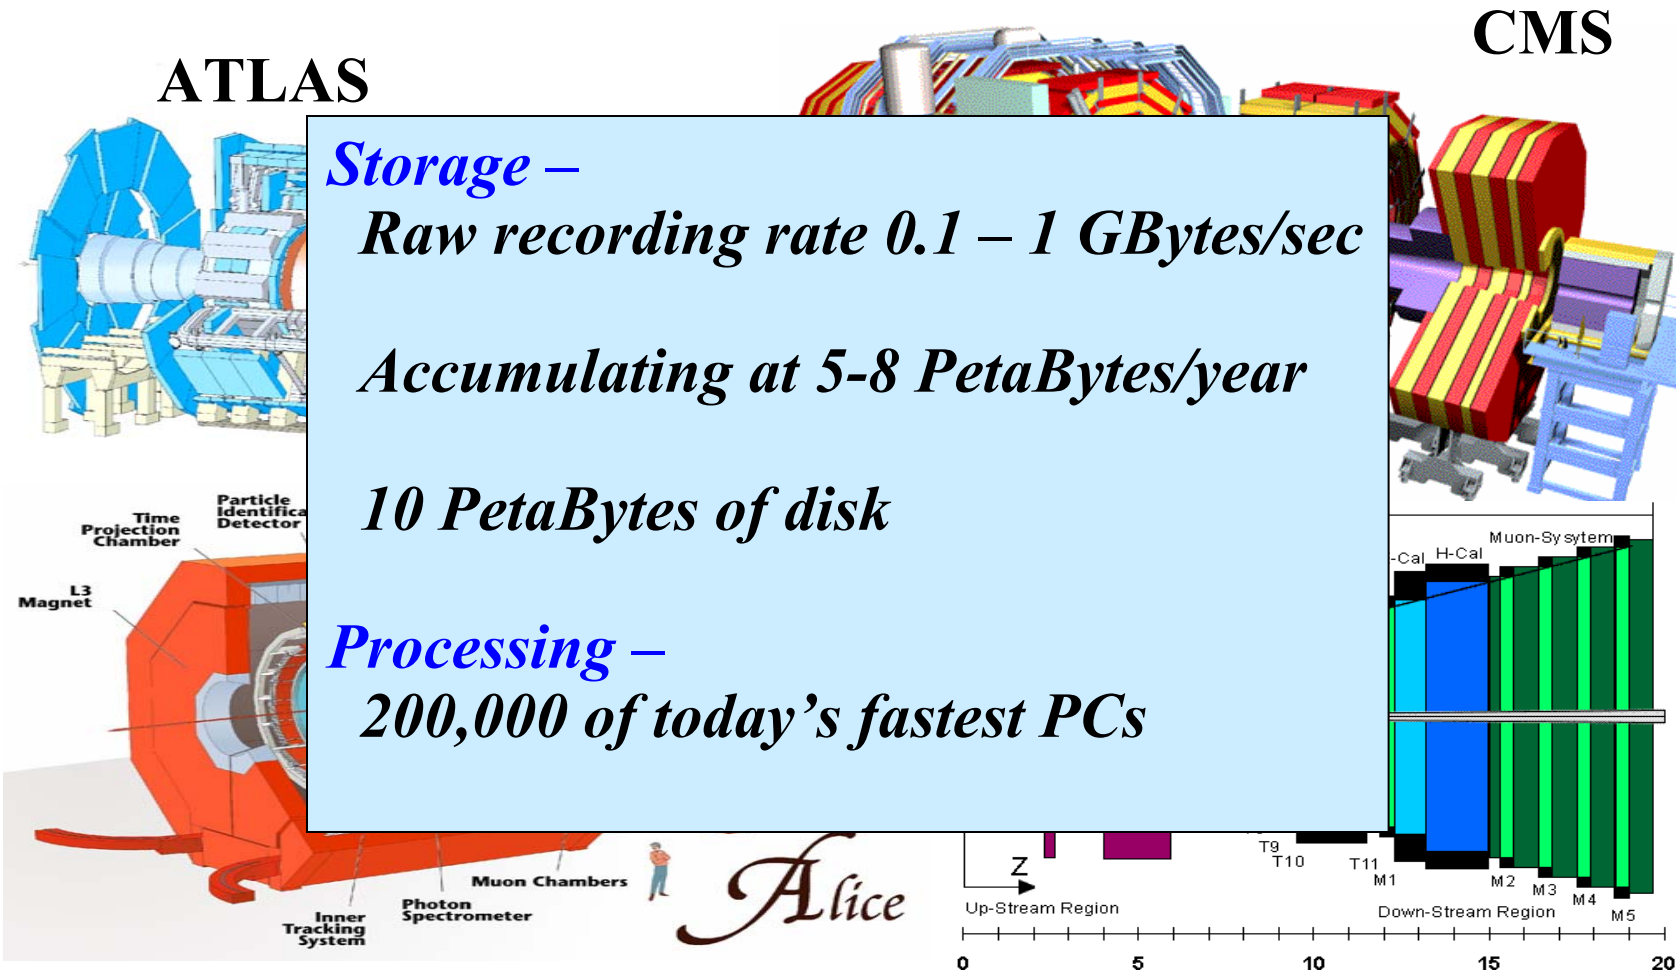

Storage –

Raw recording rate 0.1 – 1 GBytes/sec

Accumulating at 5-8 PetaBytes/year

10 PetaBytes of disk

Processing –

200,000 of today's fastest PCs

The Computing Capacity Required

Summary of Computing Capacity Required for all LHC Experiments in 2008

	CERN			Other	Total	CERN as	Total	CERN as
	Tier 0	Tier 1	Total	Tier 1	Tier 1	% of Tier 1	Tier 0 + 1	% of total Tier 0 + 1
Processing (K SI2000)	12,000	8,000	20,000	49,000	**57,000**	14%	**69,000**	29%
Disk (PetaBytes)	1.1	1.0	2.1	8.7	**9.7**	10%	**10.8**	20%
Magnetic tape (PetaBytes)	12.3	1.2	13.5	20.3	**21.6**	6%	**33.9**	40%

- CERN will provide the data reconstruction & recording service (Tier 0) -- but only a small part of the analysis capacity
- **current planning** for capacity at CERN + principal Regional Centres
  - 2002: 650 KSI2000 → <1% of capacity required in 2008
  - 2005: 6,600 KSI2000 → < 10% of 2008 capacity

Multi-Tiered View of LHC Computing

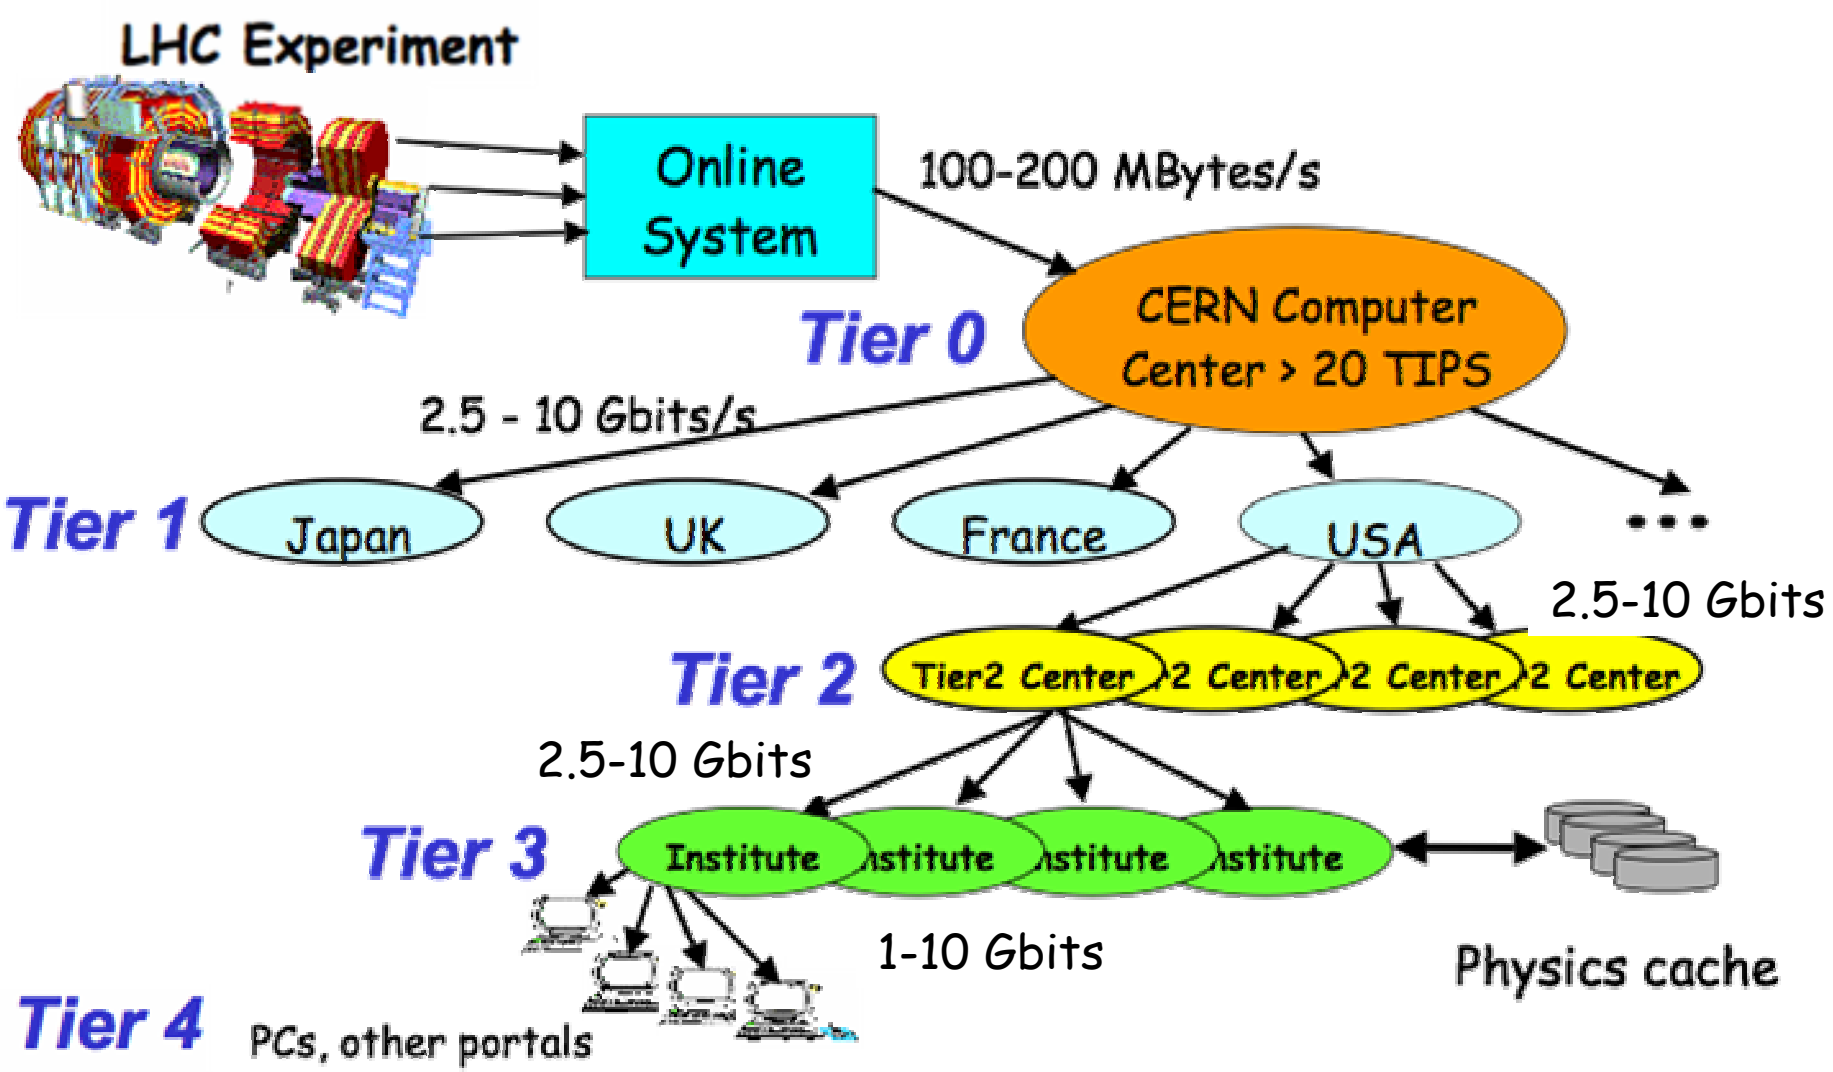

LHC Computing Model

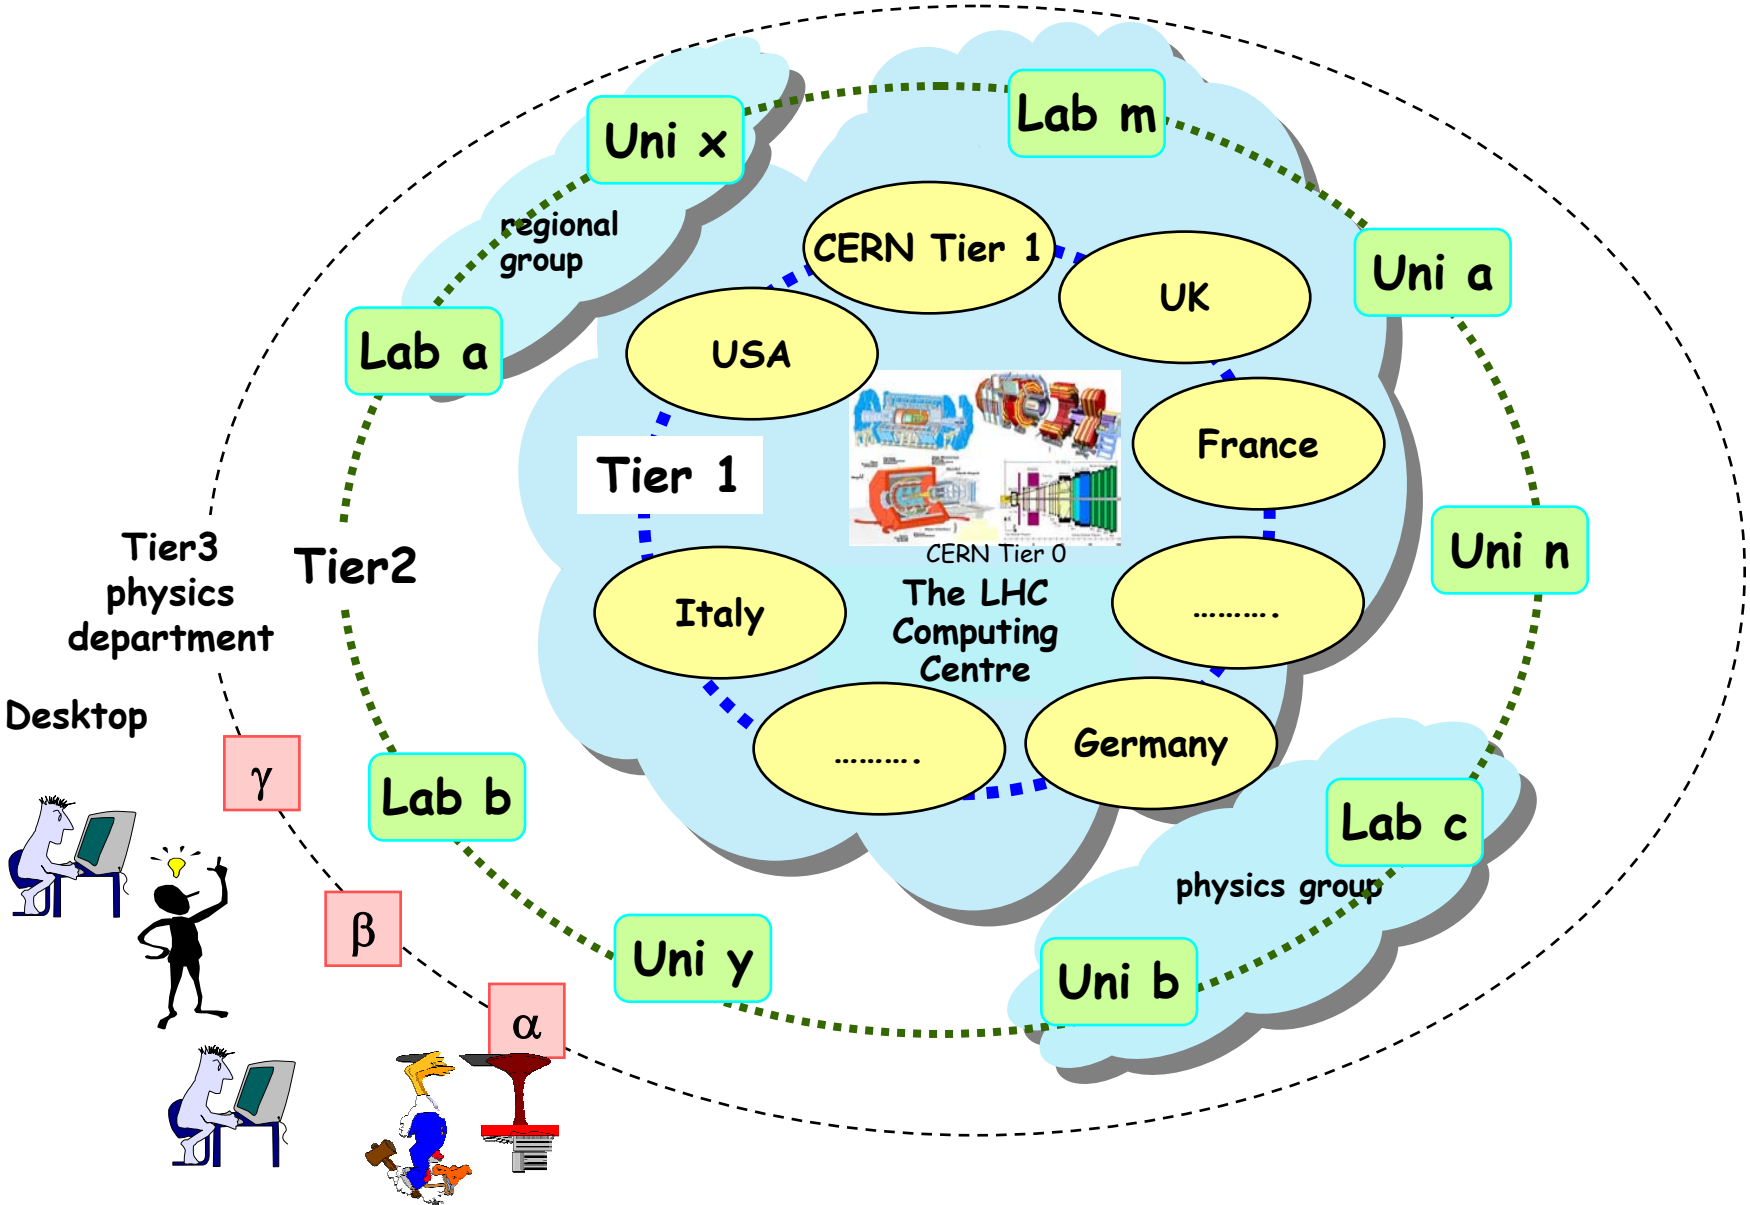

The LCG Project

The LHC Computing Grid Project

Goals:

Prepare and deploy the computing environment for the LHC experiments

- Common applications, tools, frameworks and environments,
- Move from testbed systems to real production services:
  - Operated and Supported 24x7 globally
  - Computing fabrics run as production physics services
  - Computing environment must be robust, stable, predictable, and supportable
- Foster collaboration, coherence of the LHC computing centres
- LCG is not a middleware development or grid technology project:

  It is a grid deployment project

The LHC Computing Grid Project

Two phases

Phase 1 – 2002-05

- Development and prototyping
- Approved by CERN Council 20 September 2001

Phase 2 - 2006-08

- Installation and operation of the full world-wide initial production Grid
- Costs (materials + staff) included in the LHC cost to completion estimates

The LHC Computing Grid Project

Phase 1 Goals –

- Prepare the LHC computing environment
  - provide the common tools and infrastructure for the physics application software
  - establish the technology for fabric, network and grid management (buy, borrow, or build)
  - develop models for building the Phase 2 Grid
  - validate the technology and models by building progressively more complex Grid prototypes
  - operate a series of data challenges for the experiments
  - maintain reasonable opportunities for the re-use of the results of the project in other fields
- Deploy a 50% model* production GRID including the committed LHC Regional Centres
- Produce a Technical Design Report for the full LHC Computing Grid to be built in Phase 2 of the project

* 50% of the complexity of one of the LHC experiments

LCG Deployment Plan
Level 1 Milestones

M1.1 - July 03 (Sept 03)	First Global Grid Service (LCG-1) available
M1.2 - June 03	Hybrid Event Store (Persistency Framework) available for general users
M1.3a - November 03	LCG-1 reliability and performance targets achieved
M1.3b - November 03	Distributed batch production using grid services
M1.4 - May 04	Distributed end-user interactive analysis from "Tier 3" centre
M1.5 - December 04	"50% prototype" (LCG-3) available
M1.6 - March 05	Full Persistency Framework
M1.7 - June 05	LHC Global Grid TDR

Schedule – Aggressive?

- To be ready for data taking in Spring 2007
- Need 1 year to procure, build and test the full LHC computing fabrics
- The Computing TDR must be written in mid-2005
- Need at least 1 year of experience in operating a production grid to validate the computing model
- Thus LCG must be running the experiments' data challenges in 2004
  - With a reasonable level of "production" service

Centres taking part in the LCG prototype service (2003-05)

around the world → around the clock

Centres taking part in the LCG prototype service – 2003-05

Tier 0

- CERN

Tier 1 Centres

- Brookhaven National Lab
- CNAF Bologna
- Fermilab
- FZK Karlsruhe
- IN2P3 Lyon
- Rutherford Appleton Lab (UK)
- University of Tokyo
- CERN

Other Centres

- Academica Sinica (Taipei)
- Barcelona
- Caltech
- GSI Darmstadt
- Italian Tier 2s(Torino, Milano, Legnaro)
- Manno (Switzerland)
- Moscow State University
- NCP National Centre for Physics (Islamabad,PK)
- NIKHEF Amsterdam
- Ohio Supercomputing Centre
- Sweden (NorduGrid)
- Tata Institute (India)
- Triumf (Canada)
- UCSD
- UK Tier 2s
- University of Florida– Gainesville
- University of Prague
- ……

Confirmed Resources: http://cern.ch/lcg/peb/rc_resources

LCG Resource Commitments – 1Q04

	CPU (kSI2K)	Disk TB	Support FTE	Tape TB
CERN	700	160	10.0	1000
Czech Republic	60	5	2.5	5
France	420	81	10.2	540
Germany	207	40	9.0	62
Holland	124	3	4.0	12
Italy	507	60	16.0	100
Japan	220	45	5.0	100
Poland	86	9	5.0	28
Russia	120	30	10.0	40
Taiwan	220	30	4.0	120
Spain	150	30	4.0	100
Sweden	179	40	2.0	40
Switzerland	26	5	2.0	40
UK	1780	455	24.0	300
USA	801	176	15.5	1741
Total	5600	1169	123.2	4228

LCG Project Implementation

Four work areas –

- Applications
- Grid Technology
- Fabrics
- Grid deployment

Applications Area

- Base support for the development process, infrastructure, tools, libraries
- Frameworks for simulation and analysis
- Projects common to several experiments
  - everything that is not an experiment-specific component is a potential candidate for a common project
  - long term advantages in use of resources, support, maintenance
- Object persistency and data management

Grid Technology in LCG

- LCG expects to obtain Grid Technology from
  - projects funded by national and regional e-science initiatives -- and
  - from industry

- concentrating ourselves on

  deploying a global grid service

A few of the Grid Projects with strong HEP collaboration

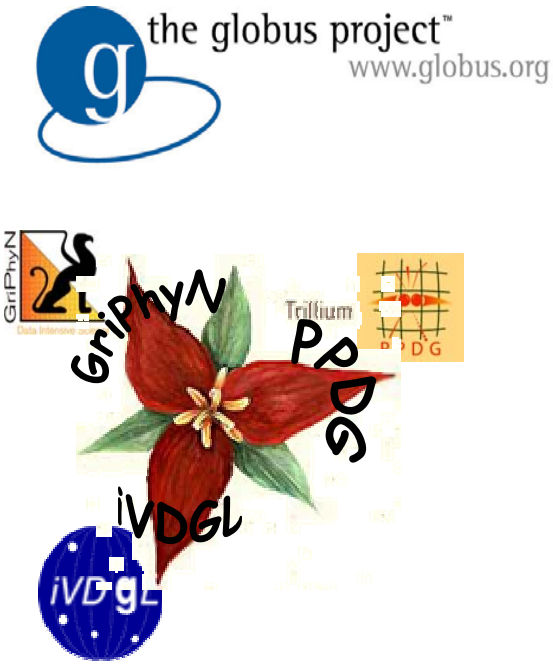

Many national, regional Grid projects -- GridPP(UK), INFN-grid(I), NorduGrid, Dutch Grid, …

US projects

European projects

Grid Technology in LCG

This area of the project is concerned with

- ensuring that the LCG requirements are known to current and potential Grid projects
- active lobbying for suitable solutions – influencing plans and priorities
- evaluating potential solutions
- negotiating support for tools developed by Grid projects

- developing a plan to supply solutions that do not emerge from other sources

- BUT this must be done with caution –
  important to avoid HEP-SPECIAL solutions
  important to migrate to standards as they emerge

  (avoid emotional attachment to prototypes)

LCG Grid Technology Organisation

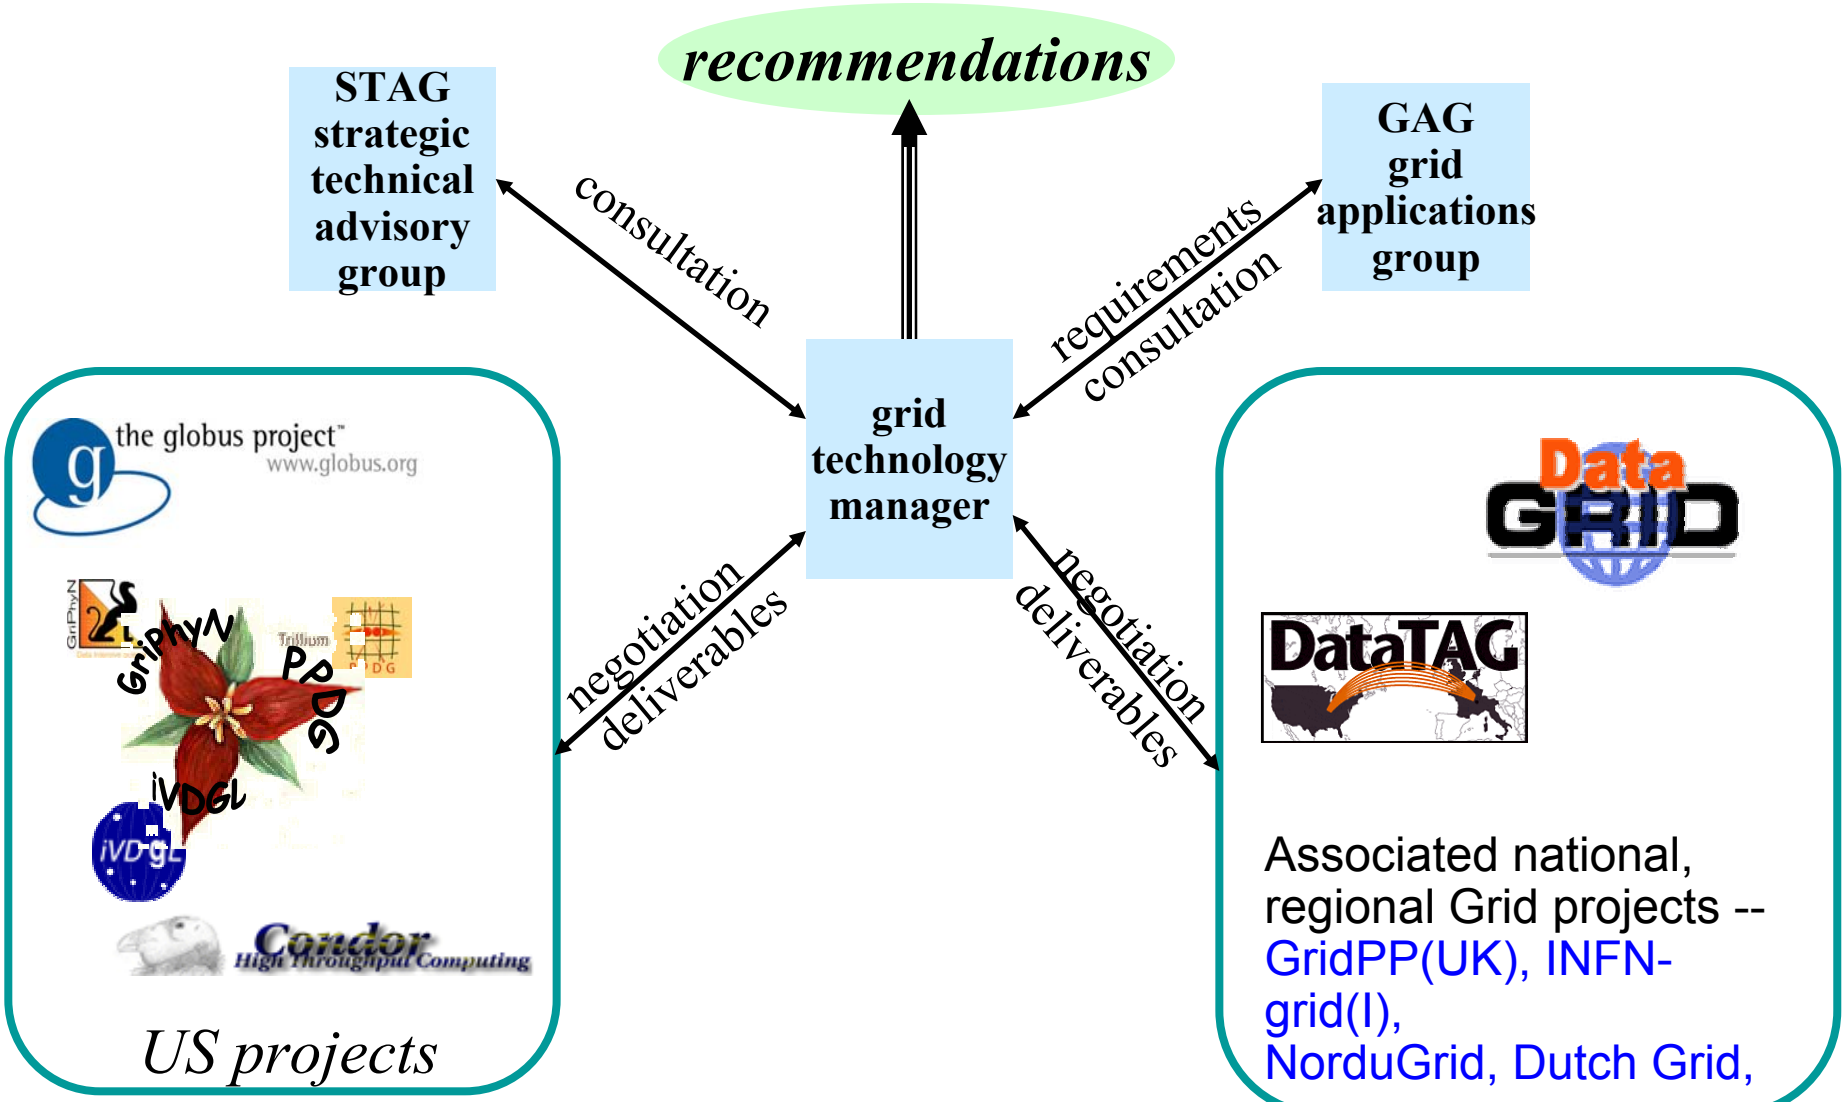

Grid Technology Status

- A base set of requirements has been defined (HEPCAL)
  - 43 use cases
  - ~2/3 of which should be satisfied ~2003 by currently funded projects
- Good experience of working with Grid projects in Europe and the United States
- Practical results from testbeds used for physics simulation campaigns

- LCG-1: (which *will* evolve)
  - VDT as the basis
  - EDG components provide higher level functionality

Grid Technology – Next Steps

- leverage the massive investments being made
  - proposals being prepared both in the EU and US
- target:
  solid (re-)engineering of current prototypes
- expect several major architectural changes before things mature

Fabric Area

- CERN Tier 0+1 centre
  - Automated systems management package – autonomic computing
  - Evolution & operation of CERN prototype – integrating the base LHC computing services into the LCG grid
- Tier 1,2 centre collaboration
  - develop/share experience on installing and operating a Grid
  - exchange information on planning and experience of large fabric management
  - look for areas for collaboration and cooperation
  - use HEPiX as the communications forum
- Technology tracking & costing
  - new technology assessment nearing completion (PASTA III)
  - re-costing of Phase II is being done in light of
    - PASTA III
    - re-assessment of experiment trigger rates, event sizes (LHCC)

Grid Deployment

Deploying a production service

Deployment Goals for LCG-1

- Production service for Data Challenges in 2H03 & 2004
  - Initially focused on batch production work
- Experience in close collaboration between the Regional Centres
  - Must have wide enough participation to understand the issues,
- Learn how to maintain and operate a global grid
- Focus on a production-quality service
  - Robustness, fault-tolerance, predictability, and supportability take precedence; additional functionality gets prioritized
- LCG should be integrated into the sites' physics computing services – should not be something apart
  - This requires coordination between participating sites in:
    - Policies and collaborative agreements
    - Resource planning and scheduling
    - Operations and Support

Timeline for the LCG computing service

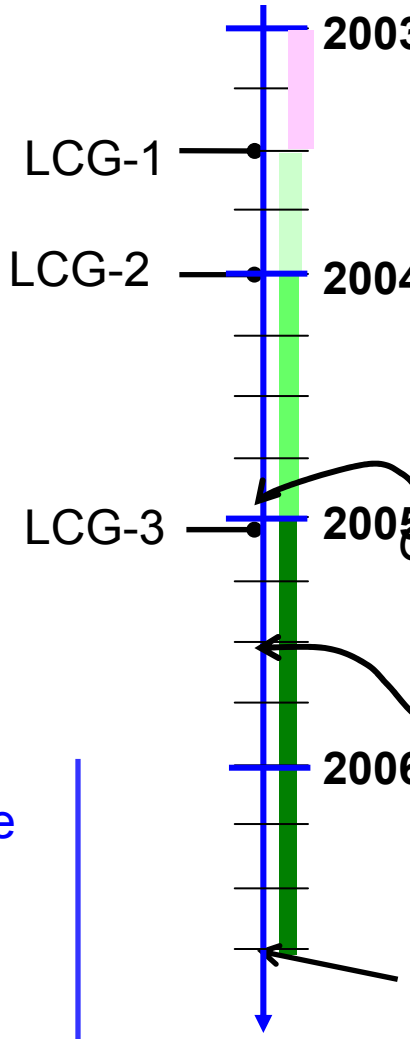

VDT, EDG tools building up to basic functionality

Stable 1st generation middleware

Developing management, operations tools

More stable 2nd generation middleware

Very stable full function middleware

Acquisition, installation, commissioning of Phase 2 service (for LHC startup)

used for simulated event productions

principal service for LHC data challenges – batch analysis and simulation

Computing model TDRs

validation of computing models

Phase 2 TDR

validation of computing service

Phase 2 service in production

The LHC Global Grid Service

LCG-1 –

First Pilot – Delivered in September 2003 –

- data replication, migration
- sustained 24 X 7 service
- including sites from three continents
- several times the capacity of the CERN facility
- and as easy to use

➢ And then evolve to the LHC production service:

- reliability, availability
- add more sites, more capacity
- service quality
- performance, efficiency
  - scheduling, data migration, data transfer
- develop interactive services
- migrate to de-facto standards as they emerge

Elements of a Production LCG Service

- Middleware:
  - Testing and certification
  - Packaging, configuration, distribution and site validation
  - Support – problem determination and resolution; feedback to middleware developers
- Operations:
  - Grid infrastructure services
  - Site fabrics run as production services
  - Operations centres – trouble and performance monitoring, problem resolution – 24x7 globally
- Support:
  - Experiment integration – ensure optimal use of system
  - User support – call centres/helpdesk – global coverage; documentation; training

General Strategy

- Use middleware, software, tools that exist
  - Developed by the various grid projects
- Integrate these tools as needed, with a well-defined testing and certification process

- Forge collaborations, common projects, agreements, to fill in the missing pieces, support, etc.
  - With grid development projects
  - With other deployment projects
  - With standards bodies (e.g. GGF)

Middleware

- Combined US and EU toolkits:

- Now:
  - VDT 1.1.8 + EDG 2.0 + GLUE schema

  - This is being used to:
  - Set up the first prototype of a production system, exercize testing & certification, deployment process, support structures
  - Address issues of integration into regional centre production environments

- This is significant – should allow inter-operation between EDG and VDT sites and LCG

Grid Deployment Organisation

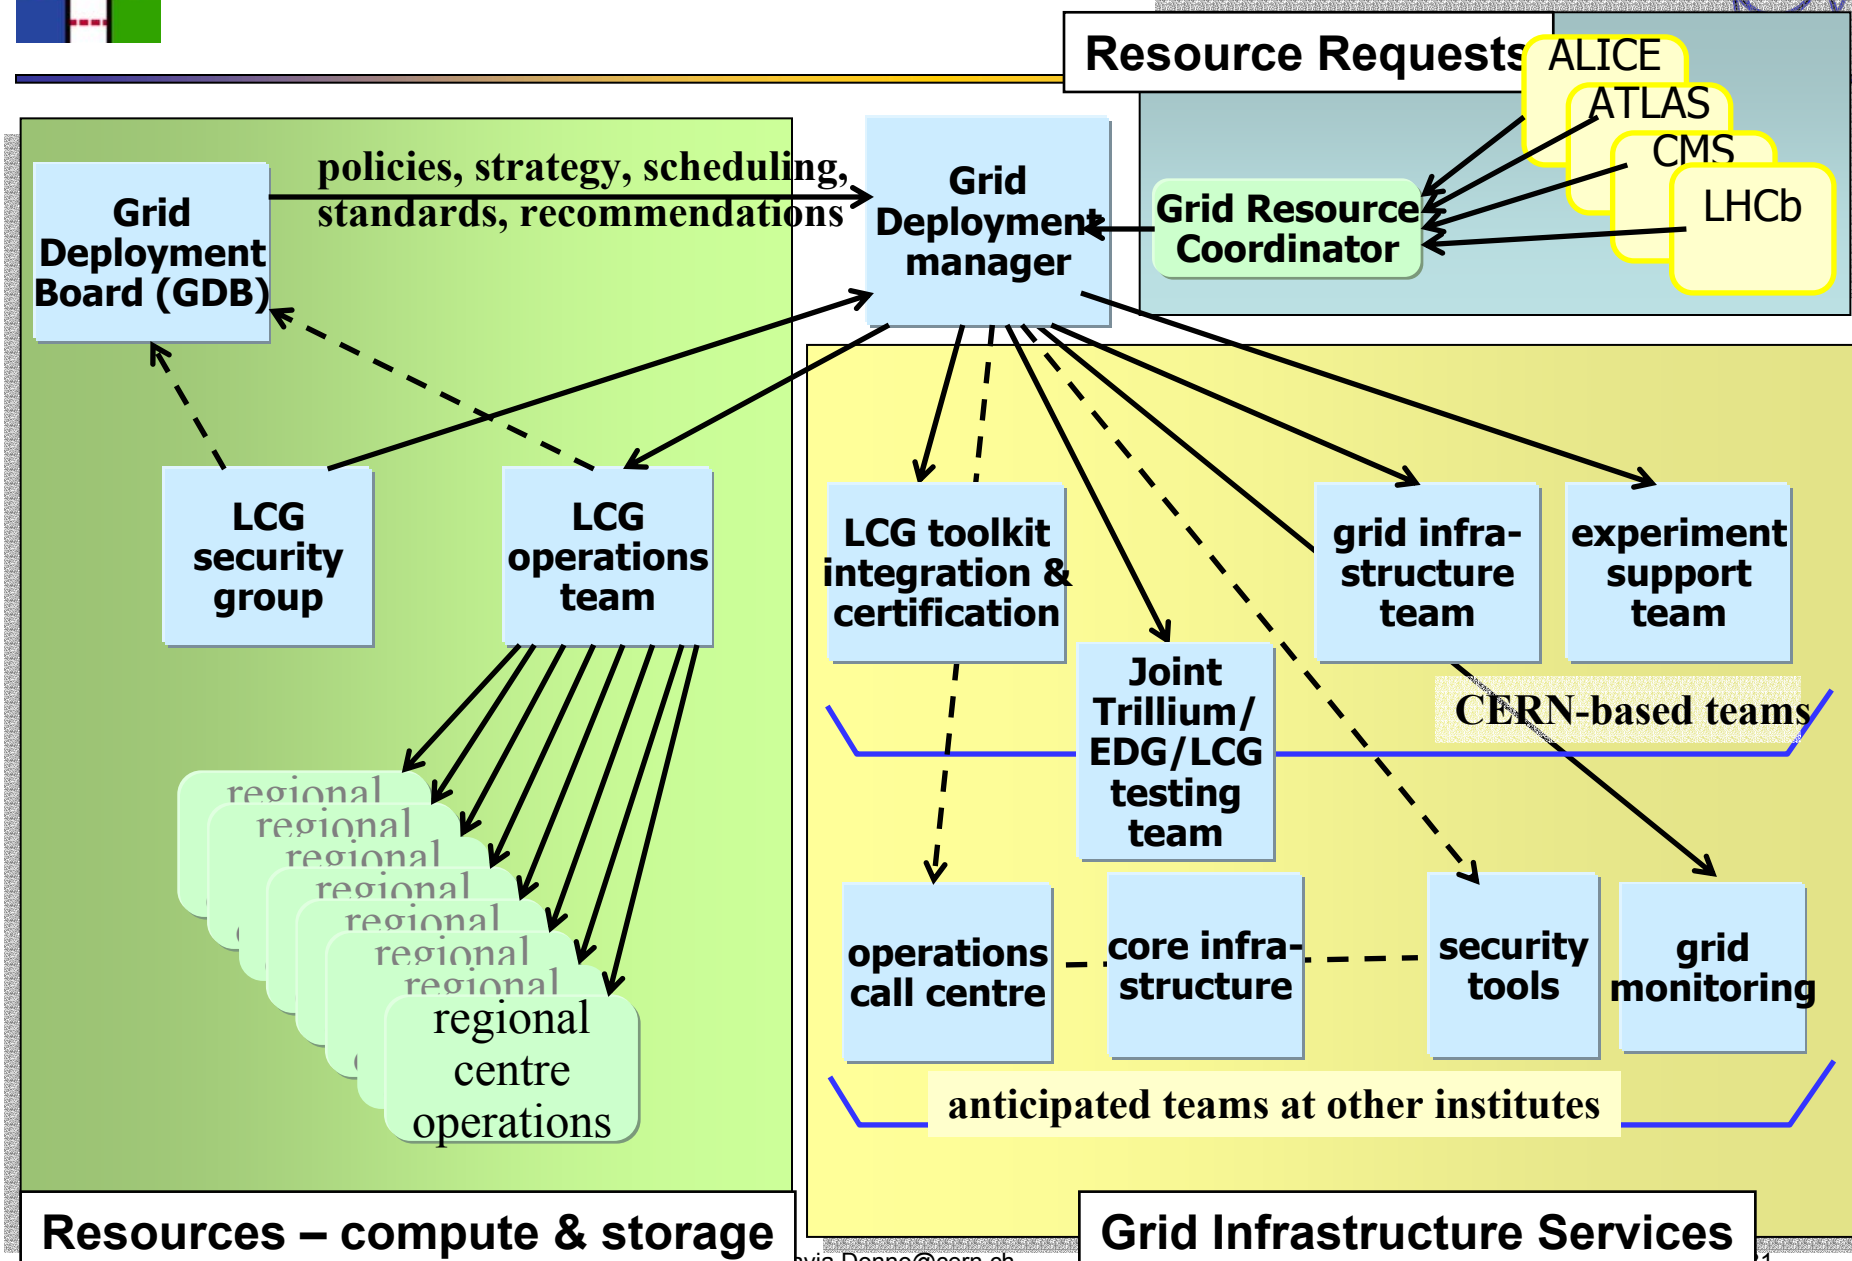

Grid Deployment Board

- Grid Deployment Board
  - representatives from the experiments and from each country with an active Regional Centre taking part in the LCG Grid Service
  - forges the agreements, takes the decisions, defines the standards and policies that are needed to set up and manage the LCG Global Grid Services
  - coordinates the planning of resources for physics and computing data challenges

- Initial task was the detailed definition of LCG-1, the initial LCG Global Grid Service
  - included defining the set of grid middleware tools to be deployed, the deployment schedule, security model, operations and support model

Certification and Testing

- Will be an ongoing major activity of LCG
  - Part of what will make LCG a production-level service
- Goals:
  - Certify/validate that middleware behaves as advertised and provides the required functionality (HEPCAL)
  - Stabilise and robustify middleware
  - Provide debugging, problem resolution *and feedback to developers*
- Testing activities at all levels
  - Component/unit tests
  - Basic functional tests, including tests of distributed (grid) services
  - Application level tests – based on HEPCAL use-cases
  - Experiment beta-testing before release
  - Site configuration verification

Certification & Testing

- Certification process – agreed a common process with EDG
  - Have agreed joint project with VDT (US):
    - VDT provide basic level (Globus, Condor) testing suites
    - We provide higher level testing
  - Will also have applications-level testing – standard benchmarks as well as experiment beta-testing, and HEPCAL tests
  - Look at using common tools and frameworks (where it makes sense) NMI/VDT-LCG
- Certification testbeds
  - Local "grid" at CERN
  - Extended to distributed test bed – U. Wisc. and others
- Site verification
  - Also an essential component
- Exception handling has not really been addressed at all …

Test and Validation process

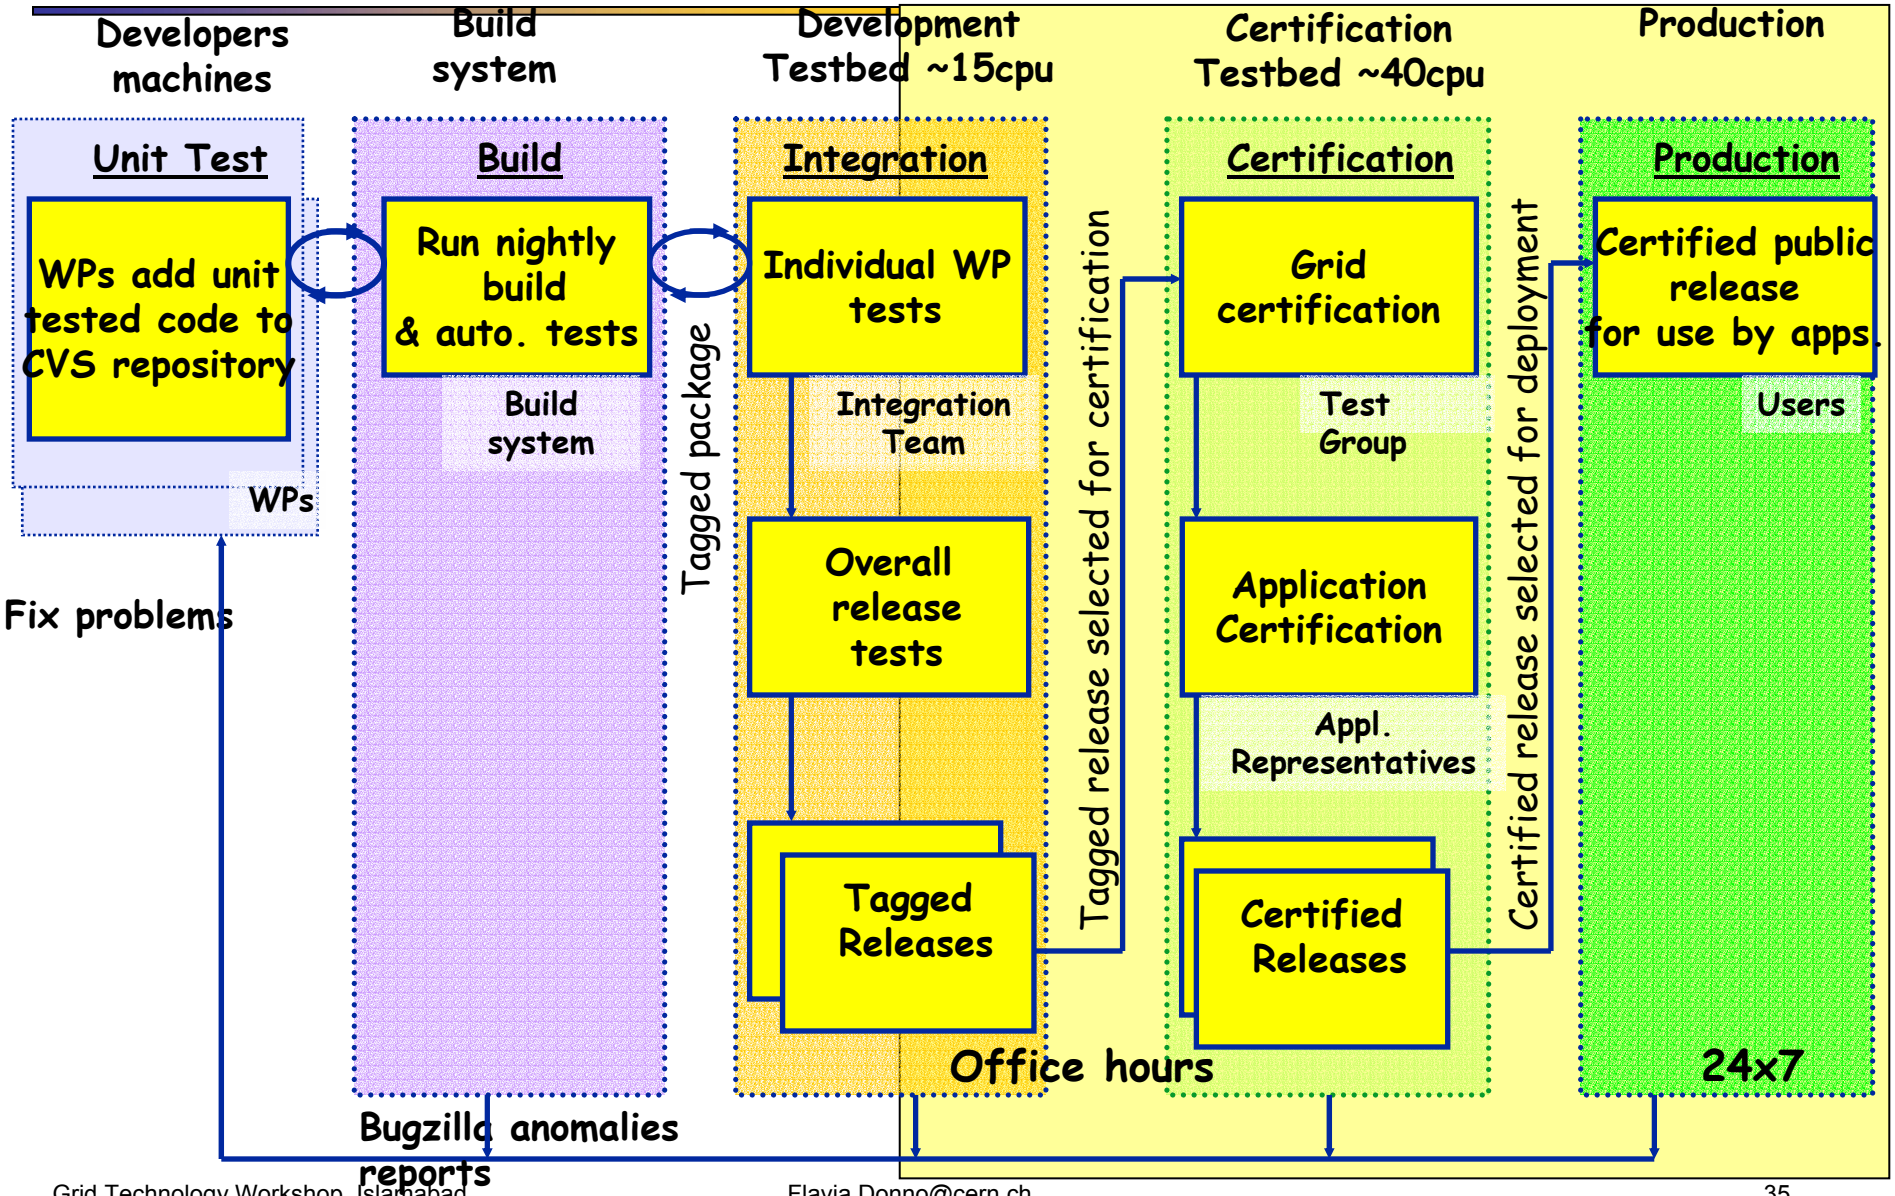

Packaging and distribution

- Obviously a major issue for a deployment project
- Want to provide a tool that satisfies needs of the participating sites,
  - Interoperate with existing tools where appropriate and necessary
  - Does not force solution on sites with established infrastructure
  - Solution for sites with nothing
- Configuration is essential component
  - Essential to understand and validate correct site configuration
  - Effort will be devoted to providing configuration tools
  - Verification of correct configuration will be required before sites join LCG
- Subject of a collaborative project

LCG Operations

- Responsible for operating and maintaining the grid infrastructure and associated services
  - Gateways, information services, resource broker etc. – i.e. grid specific services
  - Will be a coordination between teams at CERN and at Regional Centres
  - Responsible also for the VO infrastructure, Authentication and Authorisation services
  - Security operations – incident response etc.

- Build Grid Operations Centre(s)
  - Performance and problem monitoring;
  - Troubleshooting and coordination with
    - site operations,
    - user support,
    - network operations etc.
  - Accounting and reporting
  - Leverage existing experience/ideas
  - Assemble monitoring, reporting, performance, etc. tools

Monitoring tools

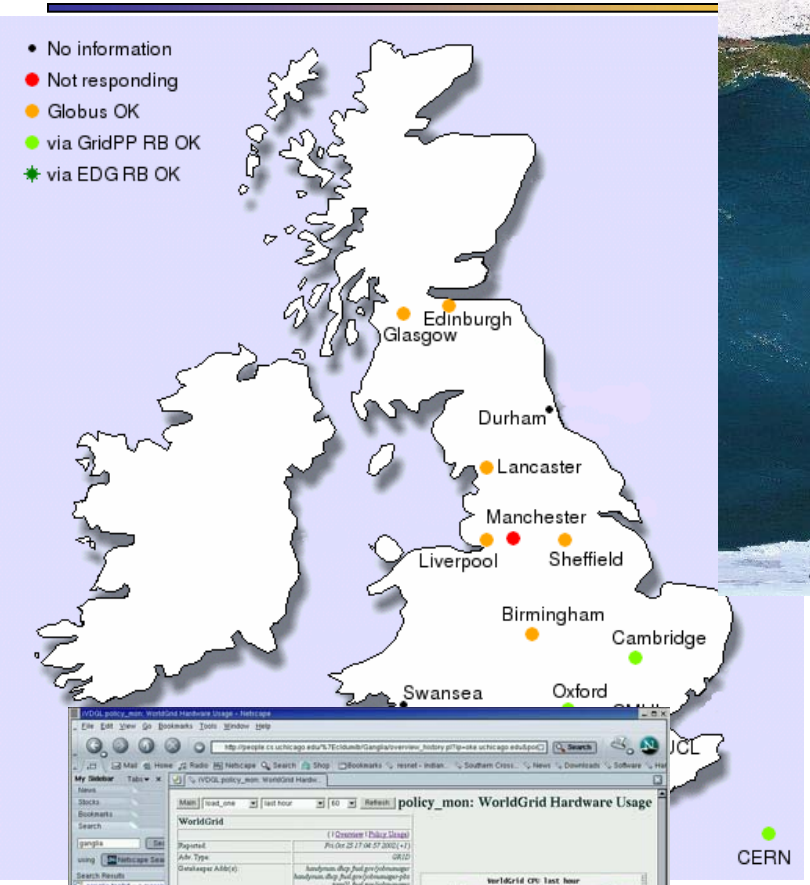

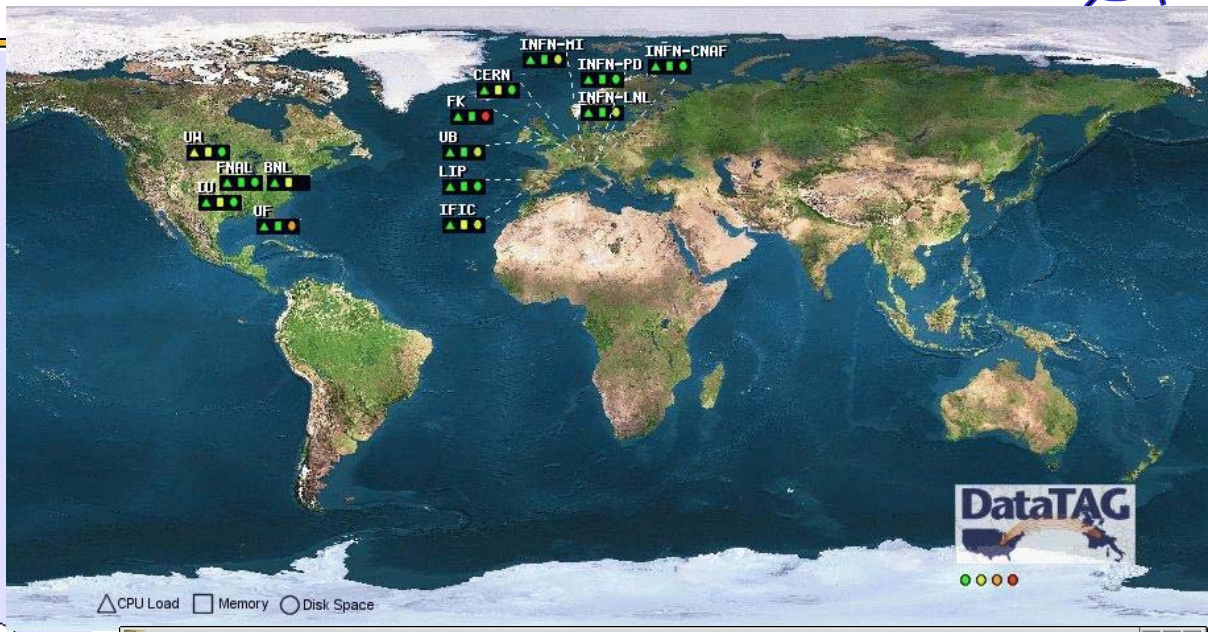

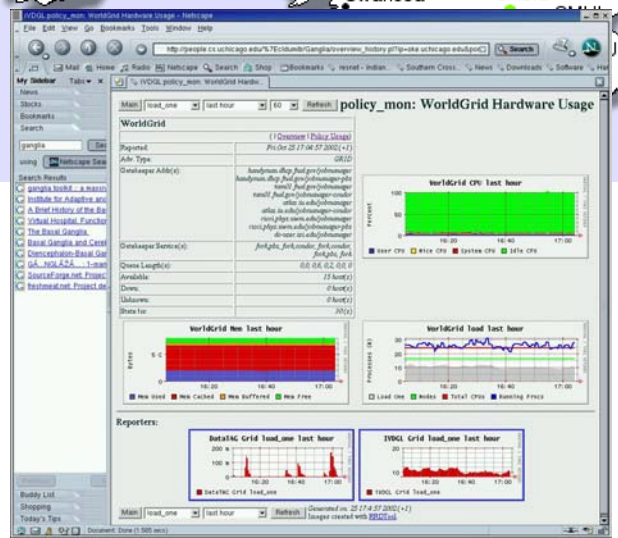

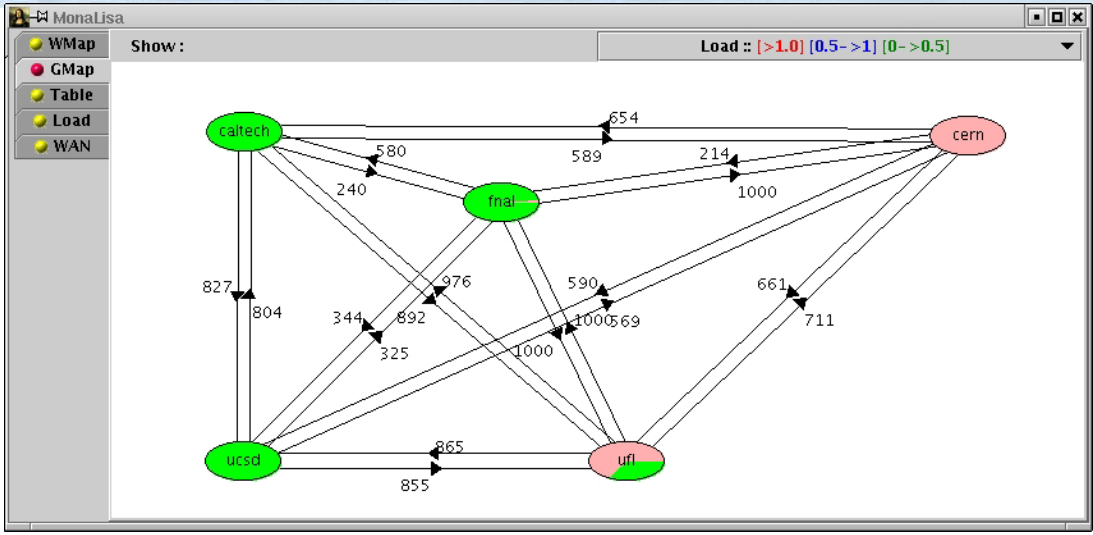

Security

- GOAL: Do not want to make exceptions for LCG services – they must run integrated into a site infrastructure, and be subject to all usual security and good management procedures and policies
- BUT: Initially, certain to need exceptions and compromises since until now most grid middleware has sidestepped security issues
- THUS: We must have a sound security policy and an agreed plan that provides for these exceptions in the short term, but shows a clear path to reach the state that the sites require
- This area represents a significant effort and must address many issues:
  - VO management
  - Usage agreements – brings up legal issues, privacy, …
  - Incident response
  - Auditing

Support Activities

- Essential for a production level service
- Experiment integration and consultancy
  - Support for data challenges
  - Ensure optimal use of resources, ensure experiment applications use middleware optimally
- Middleware support – problem determination, resolution, feedback to developers
- Call centres: 24x7 support, single point of contact
  - User support for expert users
  - Coordination of local support activities
  - Documentation
  - Training
  - Collaborate with operations centres, local user support (helpdesks)

Future Strategy

- Many LCG sites
  - Participate in other grids
  - Provide resources for other HEP experiments
  - Provide resources for other sciences
- LCG cannot exist in isolation
  - Must collaborate on standards, projects and implementations of mutual benefit
  - Essential to benefit from experience of currently running experiments trying to use grid services

Deployment Summary

- Deploy middleware to support essential functionality, but goal is to evolve and incrementally add functionality
  - Added value is to robustify, support and make into a 24x7 production service
- How?
  - Certification & test procedure – tight feedback to developers
    - must develop support agreements with grid projects to ensure this
  - Define missing functionality – require from providers
  - Provide documentation and training
  - Provide missing operational services
  - Provide a 24x7 Operations and Call Centre
    - Guarantee to respond
    - Single point of contact for a user
  - Make software easy to install – facilitate new centres joining
- Deployment is a major activity of LCG
  - Encompasses all operational and practical aspects of a grid
  - There is a lot of work already done that must be leveraged
  - Many opportunities for synergy and collaboration

Conclusions

- Moving from development to production is difficult
- Requires a lot of detailed work – needs significant investment

- There is a growing body of experience that must be built upon

- There is a good chance now to build common toolkits, share developments, and work on certification, packaging etc.

- We are forced to interoperate with other HENP experiments, other science applications: LCG cannot exist in isolation
  - This is a good thing, although it makes life harder initially

- The success of LCG will come from active collaborations with all centers participating. We are looking forward to a good collaboration with NCP.