



Enabling Grids for E-scienceE

Operating the World's largest grid infrastructure

Roberto Barbera

University of Catania and INFN

Grids@work

Sophia Antipolis, 10.10.2005


www.eu-egee.org



- **Scope and purpose of the activity**
- **Organisation**
- **Major tasks**
- **Interaction points**

24% Joint Research

- JRA1:** Middleware Engineering and Integration
- JRA2:** Quality Assurance
- JRA3:** Security
- JRA4:** Network Services Development



48% Services

- SA1:** Grid Operations, Support and Management
- SA2:** Network Resource Provision

28% Networking

- NA1:** Management
- NA2:** Dissemination and Outreach
- NA3:** User Training and Education
- NA4:** Application Identification and Support
- NA5:** Policy and International Cooperation

Emphasis in EGEE is on operating a *production Grid* and on supporting the end-users.

- Provide access to and operate a **production** grid infrastructure
 - Different user communities -> multiple Vos
 - Facilities in Europe and other collaborating sites
 - Make best use of existing grid initiatives
 - Build upon EGEE 1 experience


- What is needed to achieve this?

- 1. Core Infrastructure Services**
IS, data management, VO, (driven by Vos)
- 2. Monitoring and Control**
Performance, operational state
Initiate corrective actions
- 3. Middleware Deployment**
Integrate, certify, package middleware components
Support for new resources, setup and operation
Feedback with middleware activities in and outside of EGEE
- 4. User and resource support**
Receive problem reports
Coordinate operational problem resolution
- 5. Grid management**
Co-ordination of the implementation with the ROCs
Negotiation of SLAs
Keep in contact with the wider Grid community
Liaison, participate in standard bodies
- 6. International Collaboration**
Interoperability with large scale grids in the US and Asia-Pacific region
Seamless access for the EGEE user community to resources
- 7. Capture and provide requirements**
Relevant for operations, deployment and (some aspects of) security
Follow-up

- Scope and purpose of the activity
- **Organisation**
- Major tasks
- Interaction points with

- **Simplification**

- EGEE 1 structure
 - OMC, CICs, ROCs, RCs
- EGEE 2 (the future)
 - Operations Coordination Centre
 - Regional Operations Centres
 - Resource Centers



- **What happened to the CICs?**

- All CICs are co-located with ROCs
- Some ROCs provide CIC services

- **====> Adjust the structure to current practice**

- Basic ROCs and ROCs with CIC functions
- Easy transition, different set of services

- **Core responsibilities**
 - Middleware integration, certification, distribution packs
 - Coordinate:
 - Deployment and support
 - Grid operation and support
 - User support activity
 - Operational security activity
 - SLAs (negotiate & monitor)
 - Interoperability
 - *Non EGEE regions*
 - *ROCs more focussed on national/regional grids*
 - Act as a ROC
 - Current CIC functions (10+ RBs....)
 - ROC for RCs in non EGEE regions
 - Located at CERN

- **Support ALL sites in their region**
 - EGEE partners and friends
- **Core Responsibilities (incomplete) -----> ALL ROCs**
 - **1st line user support** (Call centre, regional training..)
 - **1st line operational support** (ROC “owns” operational problems)
 - Coordination
 - Deployment of middleware releases to its RCs
 - With national and regional grid projects
 - Regional Grid security (Incident responds teams (with RCs))
 - Negotiate resources for new VOs
 - Manage SLAs
 - Run infrastructure services
 - Support EGEE **production AND pre-production** services

- **Additional Roles**
 - Who?
 - Current CICs & ROCs with sufficient resources and expertise
 - Operations management
 - Operations Center on duty shifts
 - Monitoring, management, troubleshooting
 - *Improve, develop and run tools*
 - User support management
 - Coordinate Joint Security Policy Group (now @ RAL)
 - Run additional grid services (including VO specific)
 - Collaborate in the release process
 - Specific aspects of certification, porting, ...
 - Security vulnerability and risk analysis (NEW)
 - Coordinate (partial) code reviews, best practice,...

- **ROC concept can serve in non EGEE regions as an operation model**

- Scope and purpose of the activity
- Organisation
- **Major tasks**
- Interaction points

- Overall:
 - **Operate Production and Pre-Production Service**
 - **Some tasks implicit described with ROCs and OCC roles**
- **Middleware testing and certification**
 - Where?
 - Central coordination, some external contribution
 - Expected Results
 - Middleware distributions for production
 - Select components from **within and external** sources
 - *Negotiate support*
 - Integration and testing could be a joint activity with JRA1
 - *Testing needs to start from day 1 (sufficiently staffed)*
 - Certification
 - *Integrated system*
 - *Co-existence/Interoperability*
 - *Deployability, functionality, configuration, management of components*
 - *Extended set of OSs*
 - *Optional integration with Virtual Data Toolkit (VDT) --> ensures US interop.*

- **Testbeds**
 - Set of testbeds at CERN for rapid setup
 - Regions contribute to well defined aspects
 - Deployment tests
 - MPI support
 - Batch systems
 - Ports to different architectures

- **Middleware deployment and support**
 - OCC coordination, ROCs coordinate and support their RCs
 - Expected Results
 - Deploy agreed set to all sites
 - *Region can support supersets (but NOT subsets)*
 - Stick to agreed schedule
 - **Service Layers (new)**
 - Core services (CE, SE, Local Catalogues...)
 - *Long update cycles (1--> 2 times a year + security driven updates)*
 - *At all sites*
 - Additional Services (Central Catalogues, IS, Monitoring, RBs)
 - *Not present at all sites (mainly some ROCs)*
 - *Shorter update cycles (on demand?)*
 - Client tools on WNs
 - *Installed in user space*
 - *New version made available by a central team*
 - *VOs select preferred version*
 - Ongoing work on simplification of installation and configuration

- **Grid Operations and Support**
 - OCC & ROCs
 - Expected Results
 - Manage the grid operation
 - Has been included in the description of the ROCs and OCC's roles

- **Grid security and incident responds**
 - Security Coordination Group
 - Central coordination of incident response

 - Security Coordination Group
 - Lead by:
 - *EGEE Security Head (PEB member) +*
 - *Middleware Security Architect*
 - *Chair of the Joint Security Policy Group (SA1)*
 - *Chair of the EUGridPMA*
 - Expected Results
 - *Coordination of security related aspects of:*
 - Architecture
 - Deployment
 - Operation
 - Include standardization work

- **Grid security and incident responds**
 - Security Coordination Group
 - Central coordination of incident response
 - Central coordination of incident response
 - Coordinated at the OCC
 - ROCs coordinate the incident responds in their region
 - Requires resources at all RCs and ROCs
 - Needs a strong mandate
 - Expected Results
 - *Minimize security risks by fast responds*
 - *Ensure best practice*
 - *EGEE wide team to react on security incidents*
 - Members from ROCs/RCs

- **Support: Virtual Organizations, Applications, Users**
 - Central coordination at OCC and all ROCs

- **Expected Results**
 - User support
 - Distributed
 - Each ROC provides front-line support for local users
 - Each ROC contributes to the overall user support (experts)
 - VOs provide user support
 - *VO filters problems*
 - Existing help desks at major centres should be integrated into the support structure
 - *Filter and inject problems into the grid support*

 - User Support
 - Call centers and helpdesks
 - *ROCs*
 - Training
 - *ROCs*
 - VO support and integration
 - *NA4 with teams like the LCG-EIS*


- **We have currently not a good model for user support**
 - Some experience from LCG (can this be mapped???)
 - Needs resources from ROCs, OCC and VOs

- **Grid Management**
 - See OCC and ROCs roles
 - ROC coordinator must have a strong presence at the OCC
- **Interoperation**
 - See OCC and ROCs roles
 - ROCs focus on national/regional grids
 - OCC non EGEE regions
 - Coexistence and common policies have to be clarified
 - NA4 has to participate in the definition of “seamless”
- **Application <----> Resource Provider Coordination**
 - See OCC and ROCs roles
 - Some resources should be made available to (most) all applications
 - This could become part of the SLAs (opportunistic usage?)
 - Needs clarification
- **Application <-> RC <-> Middleware Coordination**
 - SA1 needs to be part of this
 - ROCs aggregate regional feedback
 - Coordination ?

- Scope and purpose of the activity
- Organisation
- Major tasks
- **Interaction points**

- **JRA1**
 - Integration and testing
 - Security
 - Deployment and operational requirements
 - Training
- **JRA2**
 - Work on QA metrics for operations
 - Link of QA and monitoring
- **NA4**
 - Resource negotiation
 - Security
 - Production Middleware Stack definition
 - User Support
 - Training
- **NA3**
 - Receiving and providing training (SA1 has provided significant training)
- **SA2**
 - Link between network operation center and grid operations

In collaboration with LCG



NorduGrid



Grid3/OSG

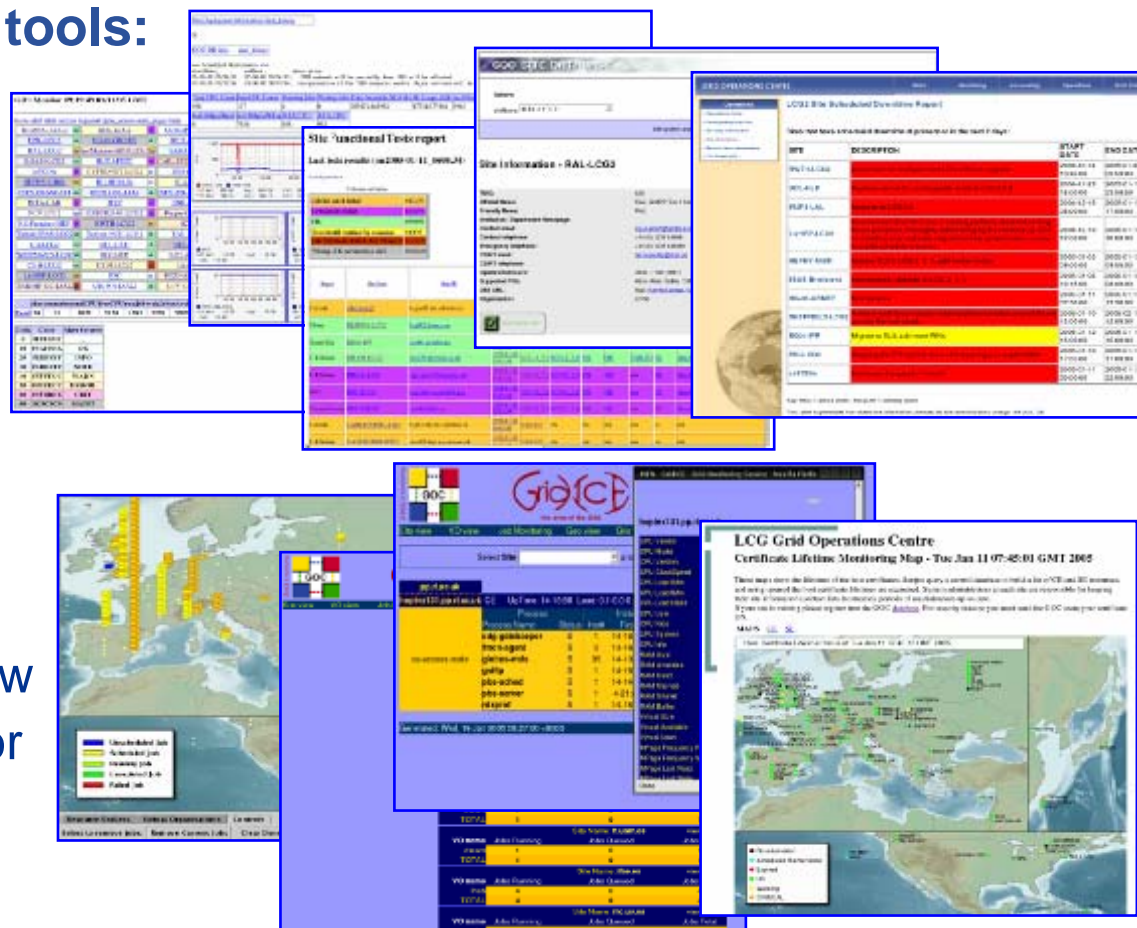


Status 25 July 2005

- Operation of Production Service: real-time display of grid operations
- Accounting Information
- Selection of Monitoring tools:

- GIS Monitor + Monitor Graphs
- Sites Functional Tests
- GOC Data Base
- Scheduled Downtimes

- Live Job Monitor
- Gridlce – VO + Fabric View
- Certificate Lifetime Monitor



- **VOs and users on the production service**

- Active VOs:

- HEP: 4 LHC, D0, CDF, Zeus, Babar
 - Biomed
 - ESR (Earth Sciences)
 - Computational chemistry
 - Magic (Astronomy)
 - EGEODE (Geo-Physics)


- Registered users in these VO: 800+

- + Many local VOs, supported by their ROCs

- **Scale of work performed:**

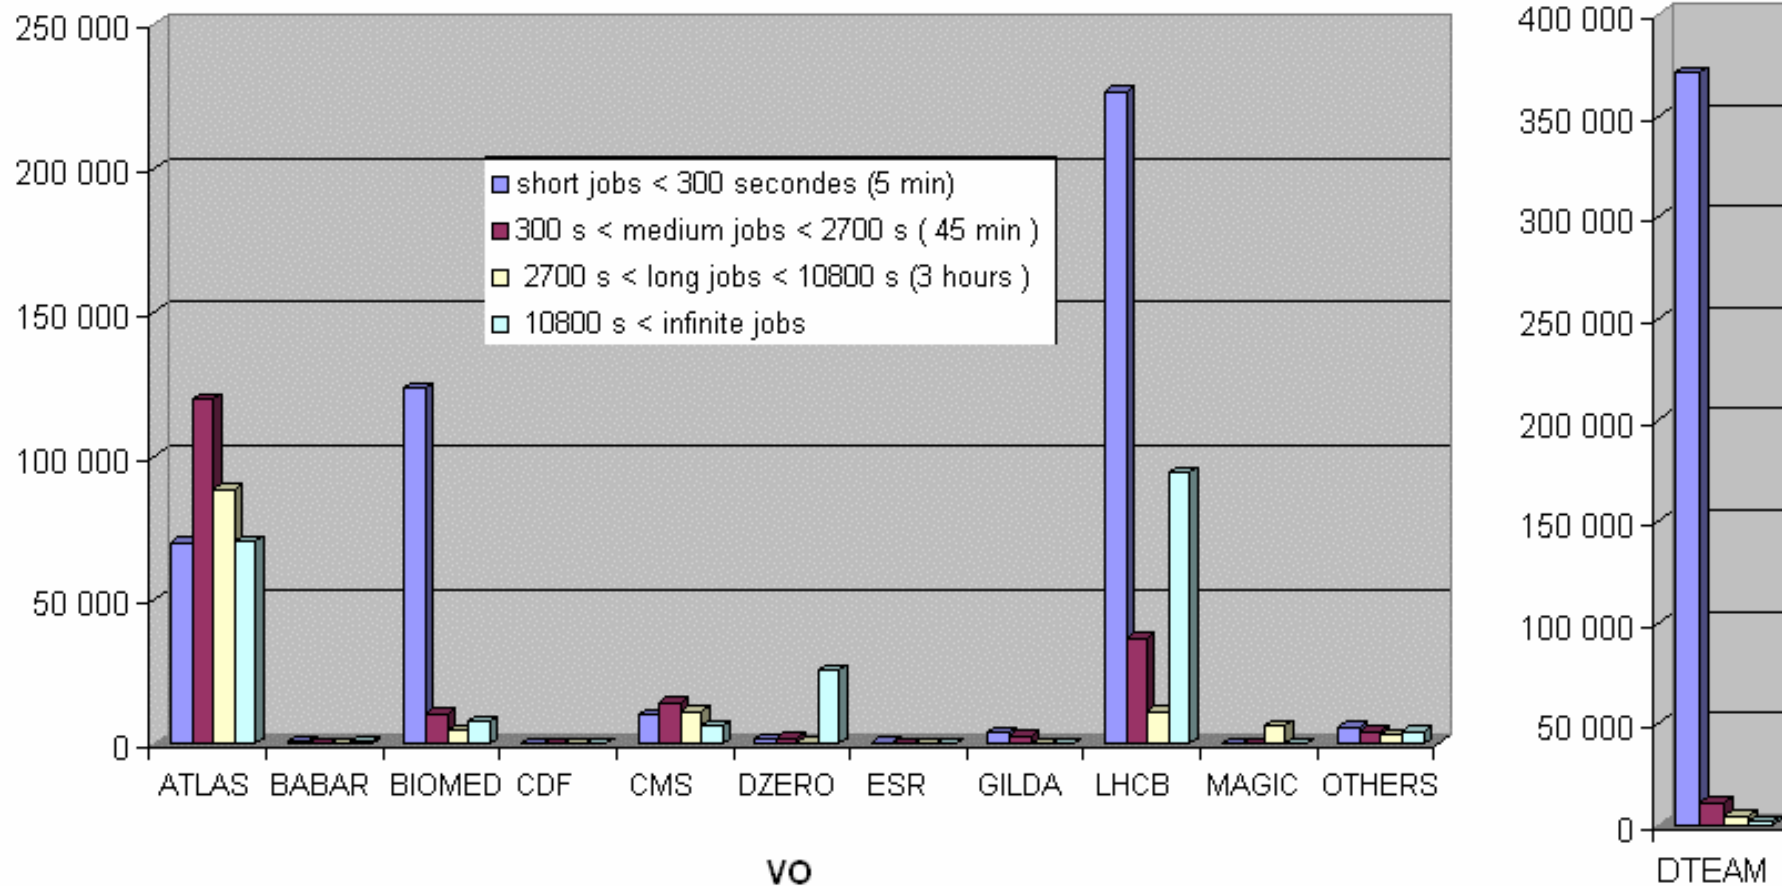
- LHC Data challenges 2004:

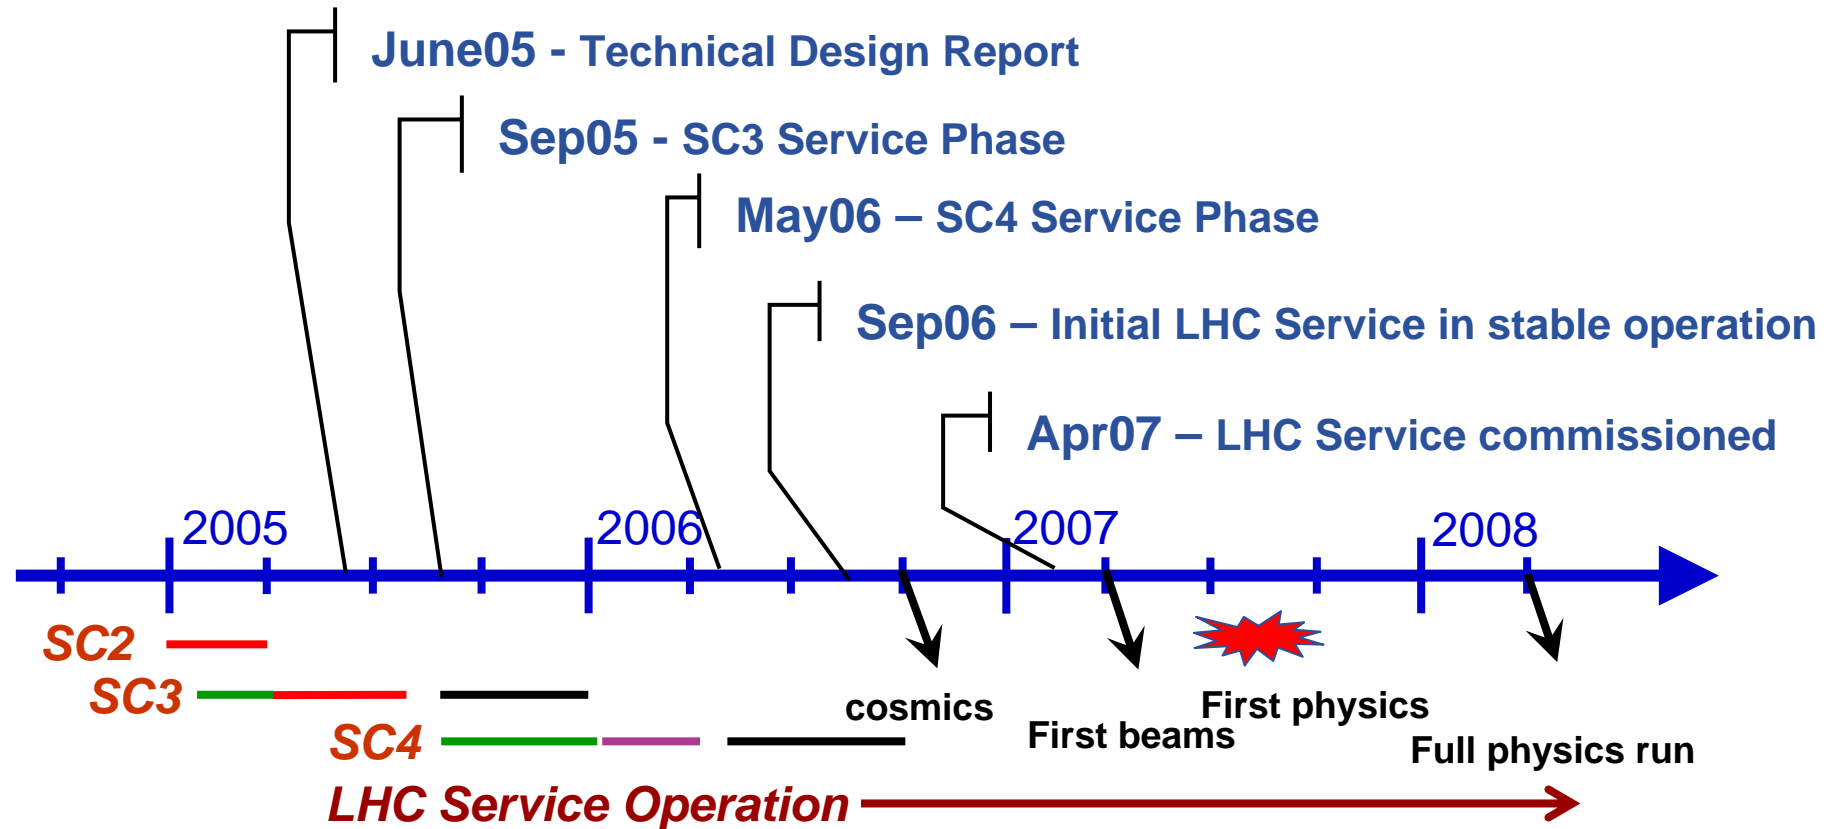
- >1 M SI2K years of CPU time (~1000 CPU years)
 - 400 TB of data generated, moved and stored
 - 1 VO achieved ~4000 simultaneous jobs (~4 times CERN grid capacity)



- Average job duration January 2005 – June 2005 for the main VOs

Number of jobs






- SC2** – Reliable data transfer (disk-network-disk) – 5 Tier-1s, aggregate 500 MB/sec sustained at CERN
- SC3** – Reliable base service – most Tier-1s, some Tier-2s – basic experiment software chain – grid data throughput 500 MB/sec, including mass storage (~25% of the nominal final throughput for the proton period)
- SC4** – All Tier-1s, major Tier-2s – capable of supporting full experiment software chain inc. analysis – sustain nominal final grid data throughput
- LHC Service in Operation** – September 2006 – ramp up to full operational capacity by April 2007 – capable of handling twice the nominal data throughput

To test Tier-0 \leftrightarrow Tier-1 \leftrightarrow Tier-2 services

- **Network service**
 - Sufficient bandwidth: ~10 Gbit/sec
 - Backup path
 - Quality of service: security, help desk, error reporting, bug fixing, ..
 - **Robust file transfer service**
 - File servers
 - File Transfer Software (GridFTP)
 - Data Management software (SRM, dCache)
 - Archiving service: tapeservers, taperobots, tapes, tapedrives, ..
 - **Sustainability**
 - Weeks in a row un-interrupted 24/7 operation
 - Manpower implications: ~7 fte/site
 - Quality of service: helpdesk, error reporting, bug fixing, ..
- Towards a stable production environment for experiments

- **>600MB/s daily average for 10 days was achieved - Midday 23rd March to Midday 2nd April**
 - Not without outages, but system showed it could recover rate again from outages
 - Load reasonable evenly divided over sites (give network bandwidth constraints of Tier-1 sites)



1) Certification test-bed

- Core functionality tested, some stress tests.
- Threshold for moving to pre-production:
 - Functionality of gLite at least that of LCG-2
 - The stability not worse than than (80)% of LCG2 on the same test-bed
 - Performance: The core functions (job submission, file reg., file lookup, file delete, data movement..) should not be less than (50)% of LCG-2

2) Pre-production

- Thresholds as for the certification test-bed.
- In addition: scalability testing.
- Applications: Overall perceived usability has to be comparable with LCG-2

3) Once thresholds achieved:

- LCG-2 is frozen,
 - except for security fixes.
 - No porting to new OS releases. This ensures that LCG-2 will be phased out with current version of OS

4) Introduction to Production:

- Major sites deploy gLite CEs in parallel with the LCG-2 CEs.
- WNs provide client libs for both stacks.
- Some of the smaller sites convert fully to gLite.
 - Incrementally, until (50)% of the resources are available through gLite.
 - Re-apply threshold tests as on pre-production (stricter?).

5) Final steps

- Migrate catalogues and data (if needed). Takes ~3 months.
- All smaller sites convert to gLite.
- Larger sites continue to provide access to LCG-2 data
 - However the LCG-2 SEs are made read-only to encourage migration of applications.

- **Keeps LCG-2 as viable fall-back**
- **Avoids having to state a drop-dead date for LCG-2 – but sets conditions**
- **Provides migration environment for applications**