Application Assessment of Production Service

Vincent Breton (CNRS)

EGEE 1st EU Review

9-11/02/2005

www.eu-egee.org

Information Society

Talk content

- Objectives of “application deployment and support” activity and its structure
- Major achievements for this past period
- Major issues and mitigation

Objectives of "application deployment and support" activity

- To identify through the dissemination partners and a well defined integration process a portfolio of early user applications from a broad range of application sectors from academia, industry and commerce

- To support development and production use of all of these applications on the EGEE infrastructure and thereby establish a strong user base on which to build a broad EGEE user community

- To initially focus on two well-defined pilot application areas, Particle Physics and Biomedicine

The role of the pilot applications – HEP and Biomedicine

- Initial area of focus to establish a strong user base on which to build a broad EGEE user community
- Provide early feedback to the infrastructure activities on their experience with application deployment and VO management
- Act as guinea pigs and provide early feedback to the middleware developers on their experience with new services

The characteristics of pilot HEP applications

- **Very large scale from project day 1**
- **Virtual Organizations were already set up at project day 1**
- **Very centralized: jobs are sent in a very organized way**
- **Multi-grid: data challenges are deployed on several grids**
  - ALICE LCG, Alien
  - ATLAS LCG, US Grid2003, Nordugrid
  - CMS LCG, US Grid2003
  - LHCb LCG, Dirac

Data Challenges – ALICE

- Phase I (Feb-May 2004)
  - 120k Pb+Pb events produced in 56k jobs
  - 1.3 million files (26TByte) in Castor@CERN
  - CPU: ~35 kSI2k years
- Phase II (Aug-Dec 2004)
  - 340k jobs
  - 40 TB produced, 200TB transferred
  - CPU: ~85 kSI2k years
- Overall ~ 15% jobs ran on LCG – will extend usage with new middleware offering improved data management

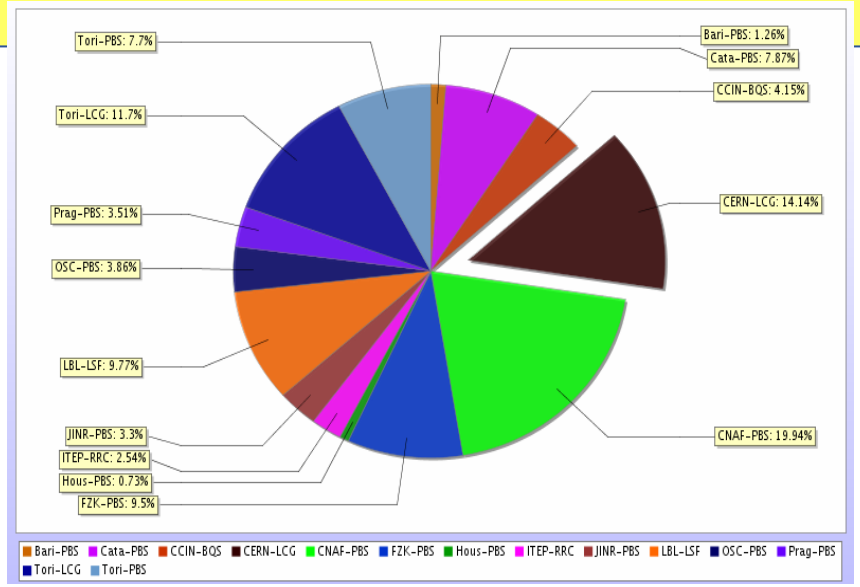

Data Challenges – CMS

March-April 2004

- ~30 M events reconstructed at Tier-0
- 25Hz reached for flow to analysis in Tier-1
  - (only once for a full day)
  - RLS, Castor, control systems, T1 storage, …

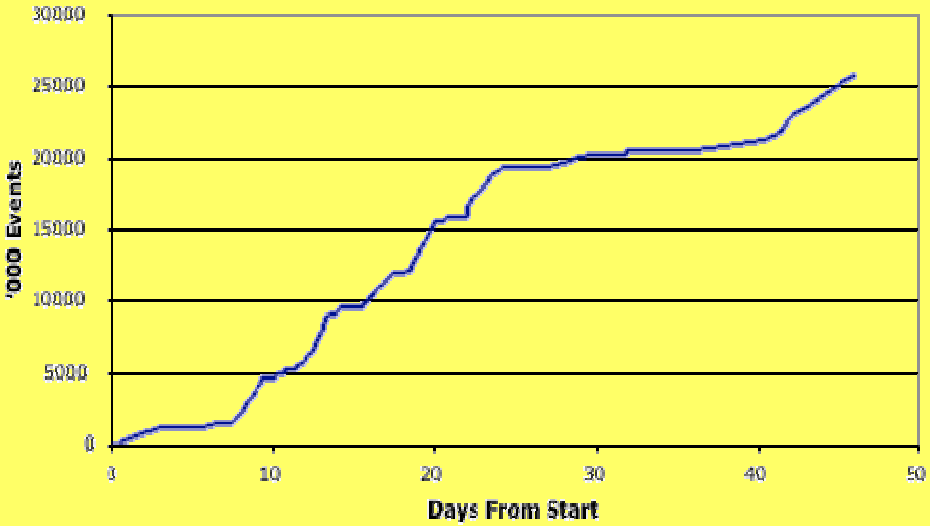

- Not a CPU challenge, but a full chain demonstration
- Pre-challenge production in 2003/04
  - **70 M Monte Carlo events (30M with Geant-4) produced**
  - **Classic and grid (CMS/LCG-0, LCG-1, Grid3) productions**

Data Challenges – LHCb

June-September 2004

- 186 M events in ~250k jobs . 61 Tbyte data
- CPU: ~370 kSI2k years (43 LCG and 20 DIRAC sites)
- Increased usage of LCG with time – 75% by September; 50% LCG overall
- Up to 3500 concurrent running jobs in LCG

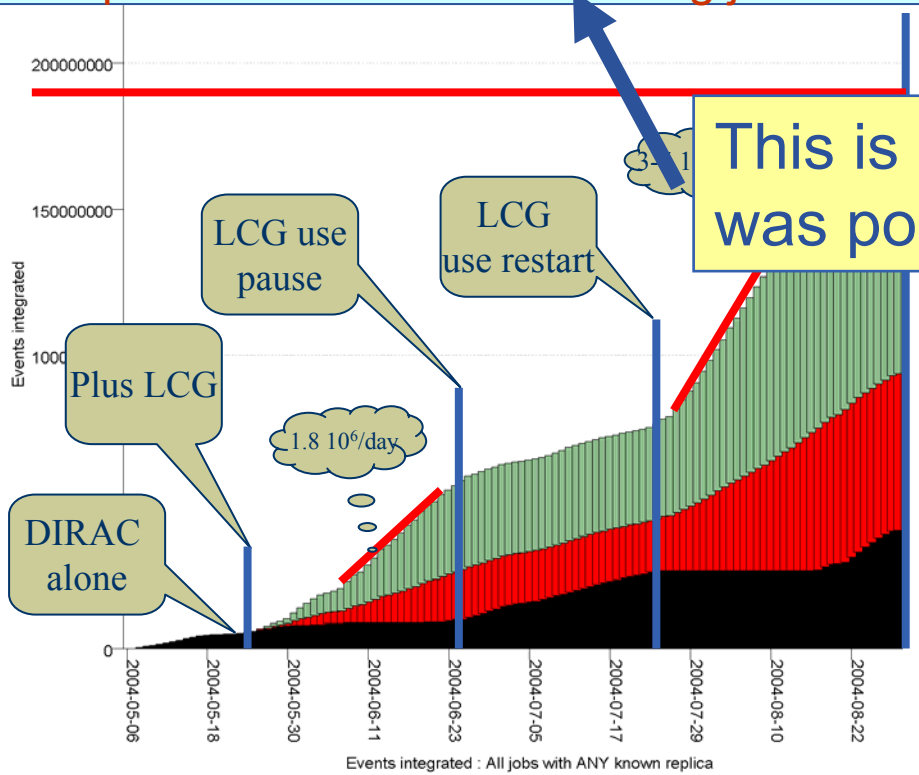

This is more than 3 times what was possible at CERN alone

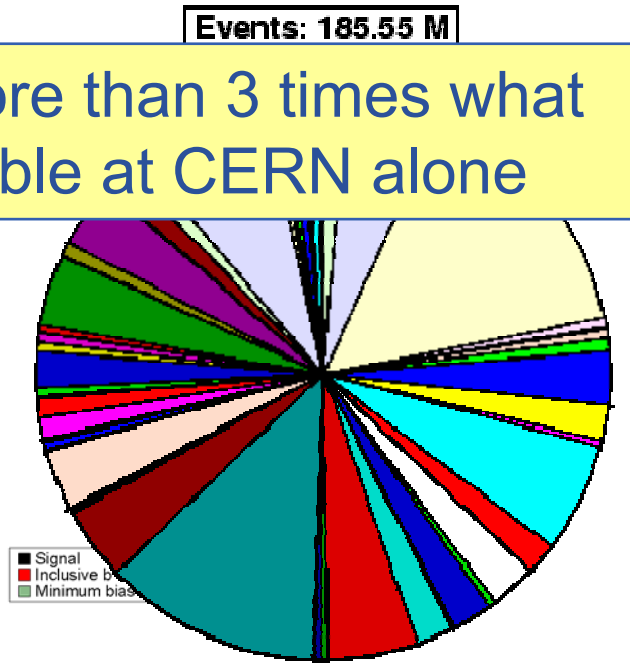

Data Challenges – ATLAS

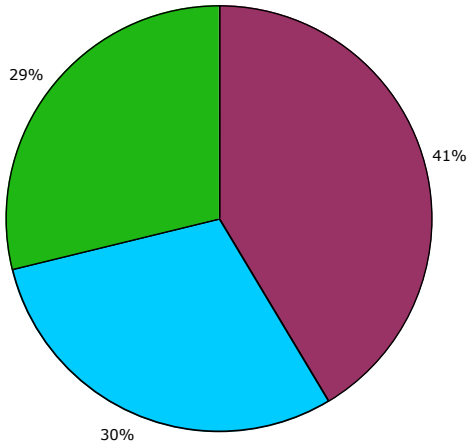

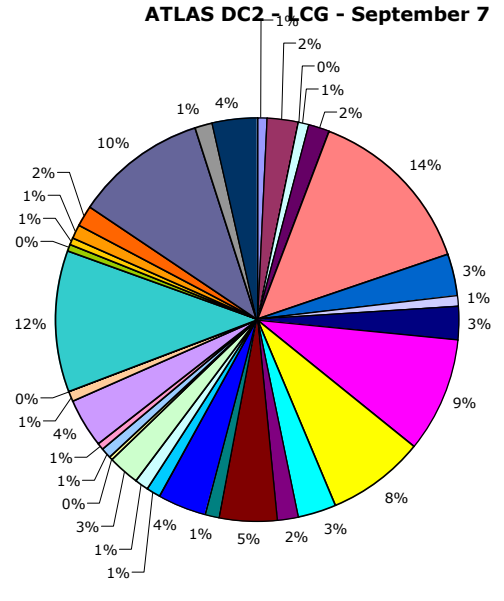

July-October 2004

- 7.7 Million events fully simulated (Geant 4) in ~ 95k jobs
- 22 Tbyte of data
- CPU: ~110 kSI2k years
- >40% produced on LCG

Overview of experiences with LHC data challenges

- **There was continual evolution throughout 2004, with LCG and experiments gaining more experience in the development and use of an expanding LCG grid**
- **All experiments had excellent relations with LCG-EIS support – a model for the future support of VOs**
- **Global job efficiencies ranged from 60-80% as experience developed – must get up to 90+% for user analysis - look to new middleware developments and tighter operational procedures**
- **Sources of problems and losses**
  - Site configuration, management and stability
  - Data Management (especially metadata handling)
  - Difficult to monitor job running and causes of failure
- **D0 in early 2005 showed that one can run with good efficiency with a set of well controlled sites**

D0 MC efficiency on LCG2 since Xmas (but small statistics)

CE	Success	Failed
bohr0001.tier2.hep.man.ac.uk	237	3
cclcgceli01.in2p3.fr	-	14
grid-ce.physik.uni-wuppertal.de	-	-
gridkap01.fzk.de	2564	19
golias25.farm.particle.cz	198	15
heplnx131.pp.rl.ac.uk	246	4
lcgce02.gridpp.rl.ac.uk	293	10
mu6.matrix.sara.nl	397	7
tbn18.nikhef.nl	154	2
Total	4089	74

- Efficiency 98 %
- System running monitored very closely by run-manager in close contact with sites

The characteristics of biomedical pilot applications

- **Prototype level at project day 1**
- **VO was created after the project kicked-off**
- **Very decentralized: application developers use the grid at their own pace**
- **Very demanding on services**
  - Compute intensive applications
  - Applications requiring large amounts of short jobs
  - Need for interactivity or guaranteed response time

- **Resources were focused on the deployment of large scale applications on LCG-2**
  - Integration of Biomed VO used to identify issues relevant to all VOs to be deployed during EGEE lifetime
  - Decentralized usage of the infrastructure highlights different weaknesses from the more centralized HEP data challenges

Status of Biomedical VO

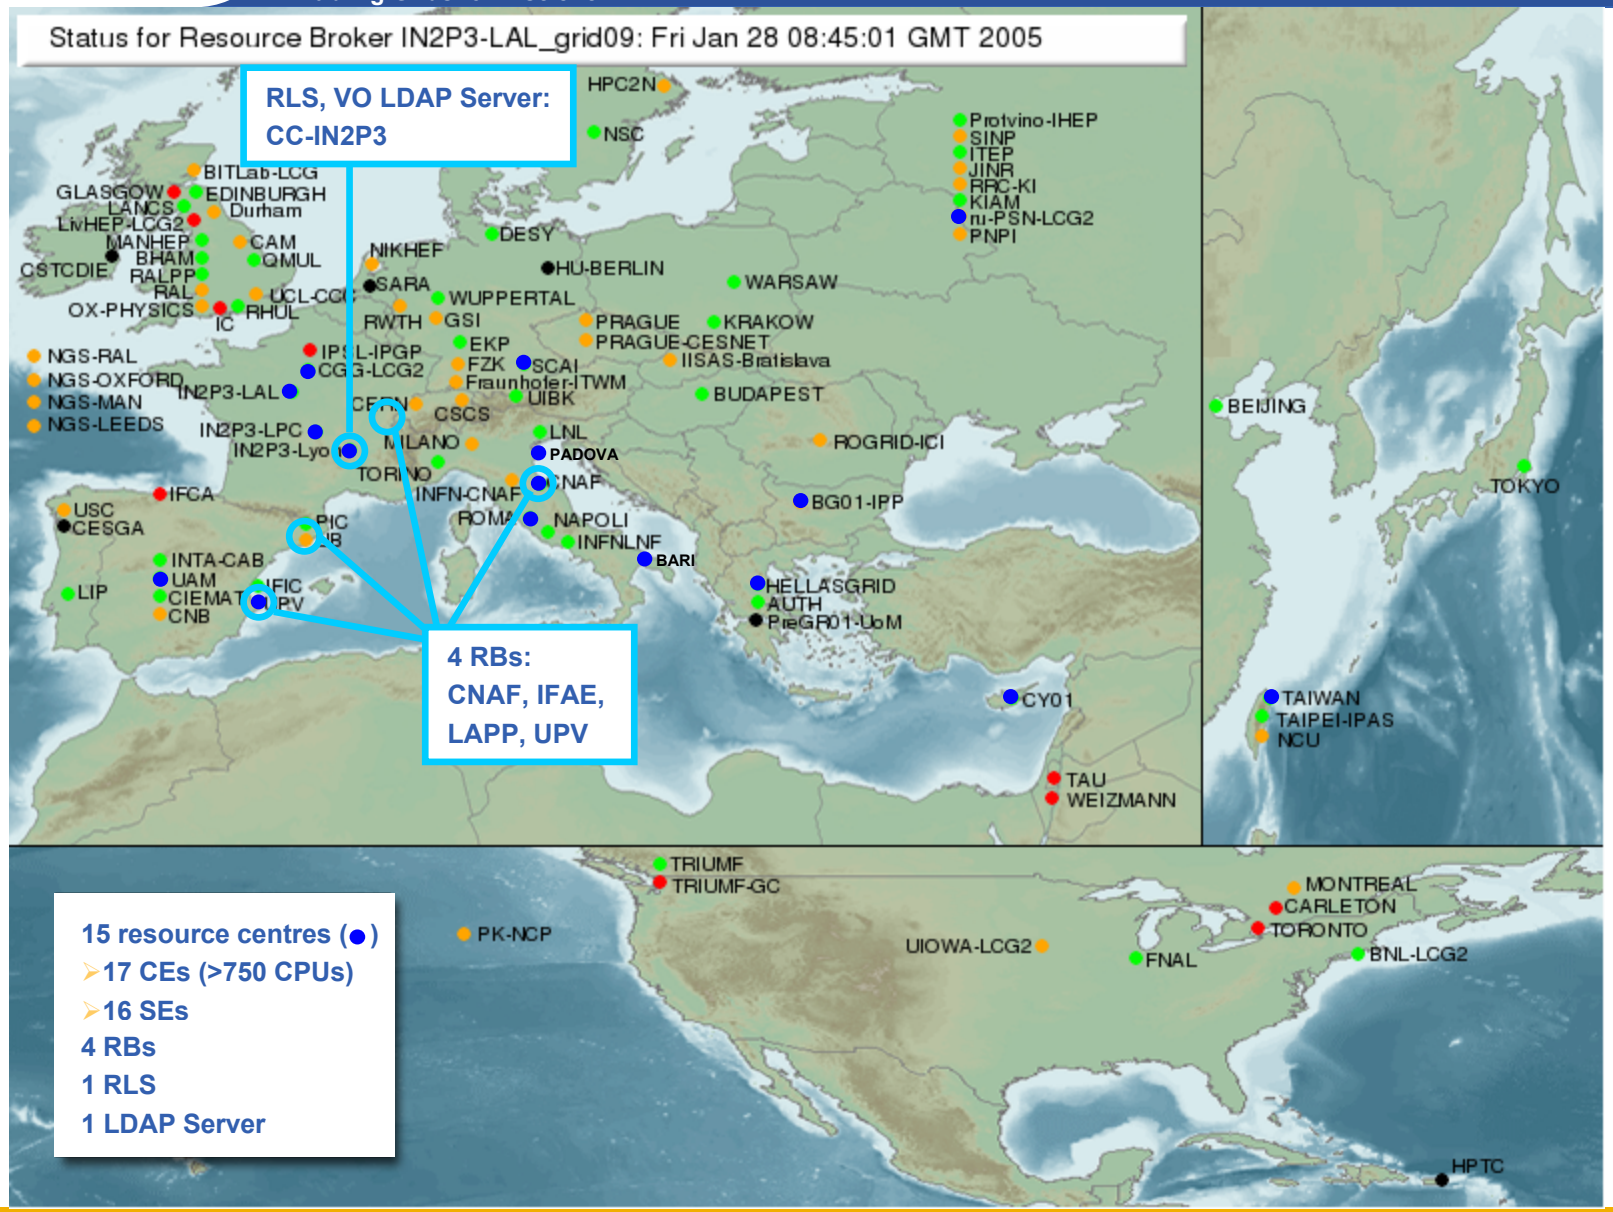

Biomedical VO: production jobs on EGEE

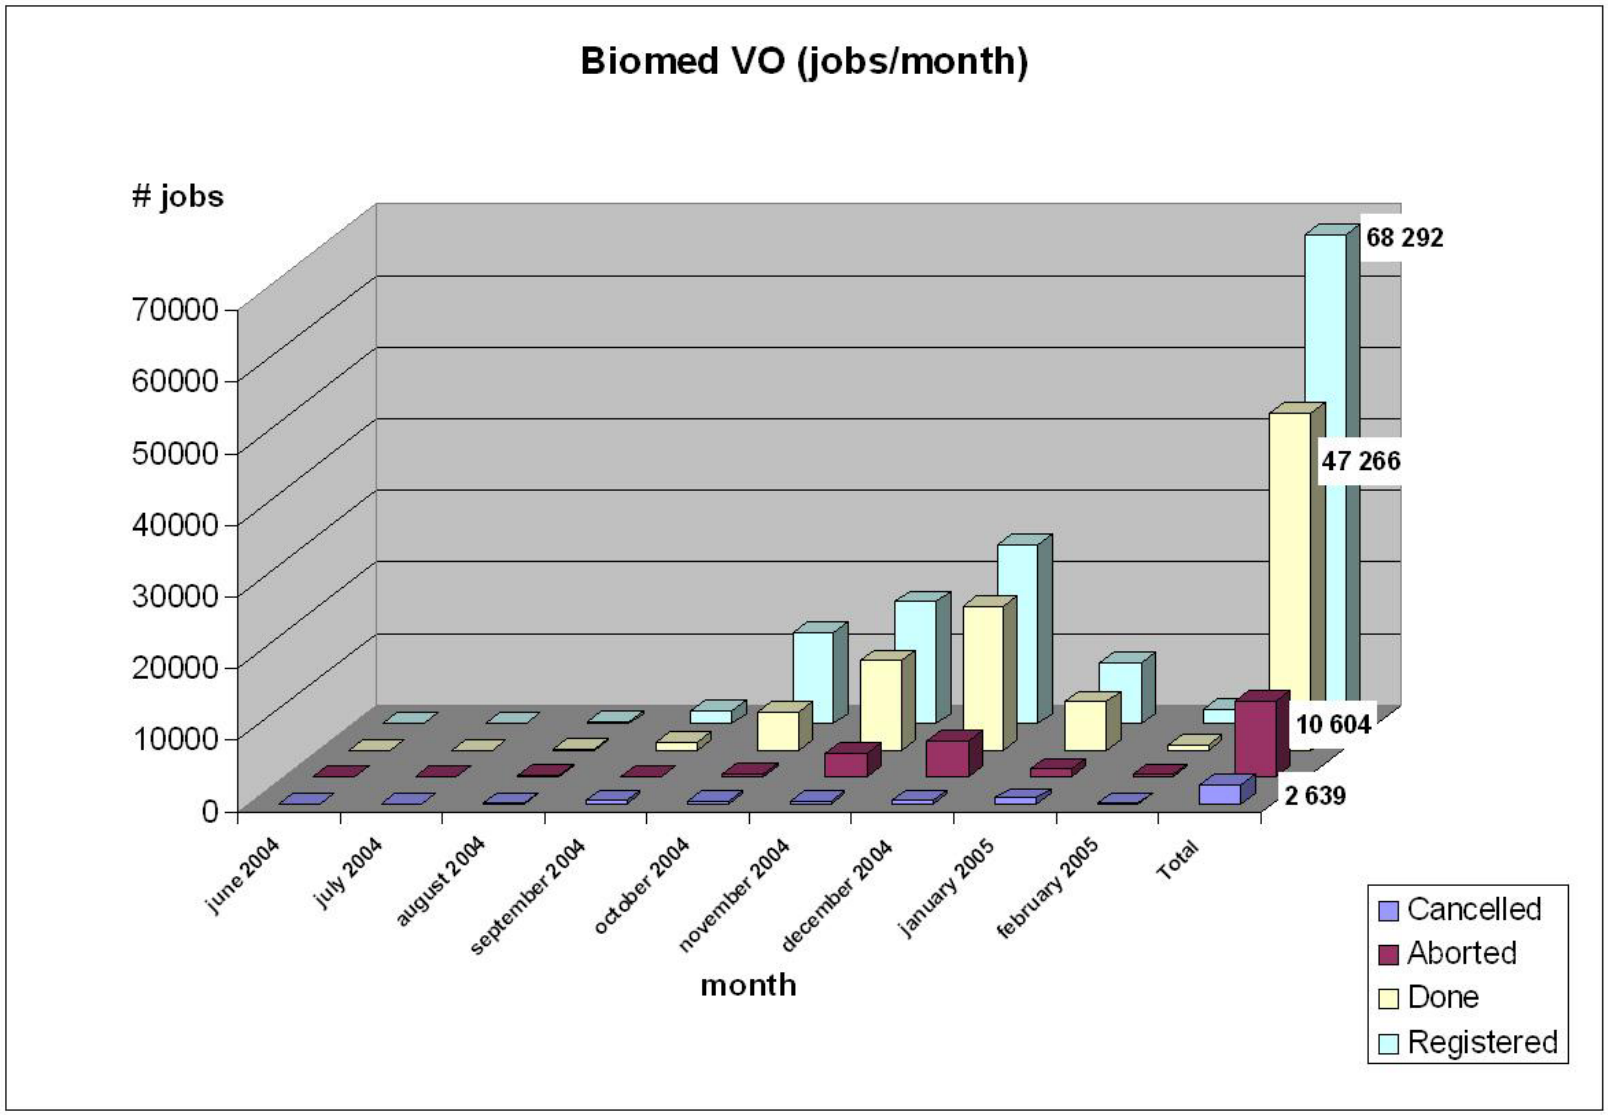

Biomedical VO: production jobs on EGEE

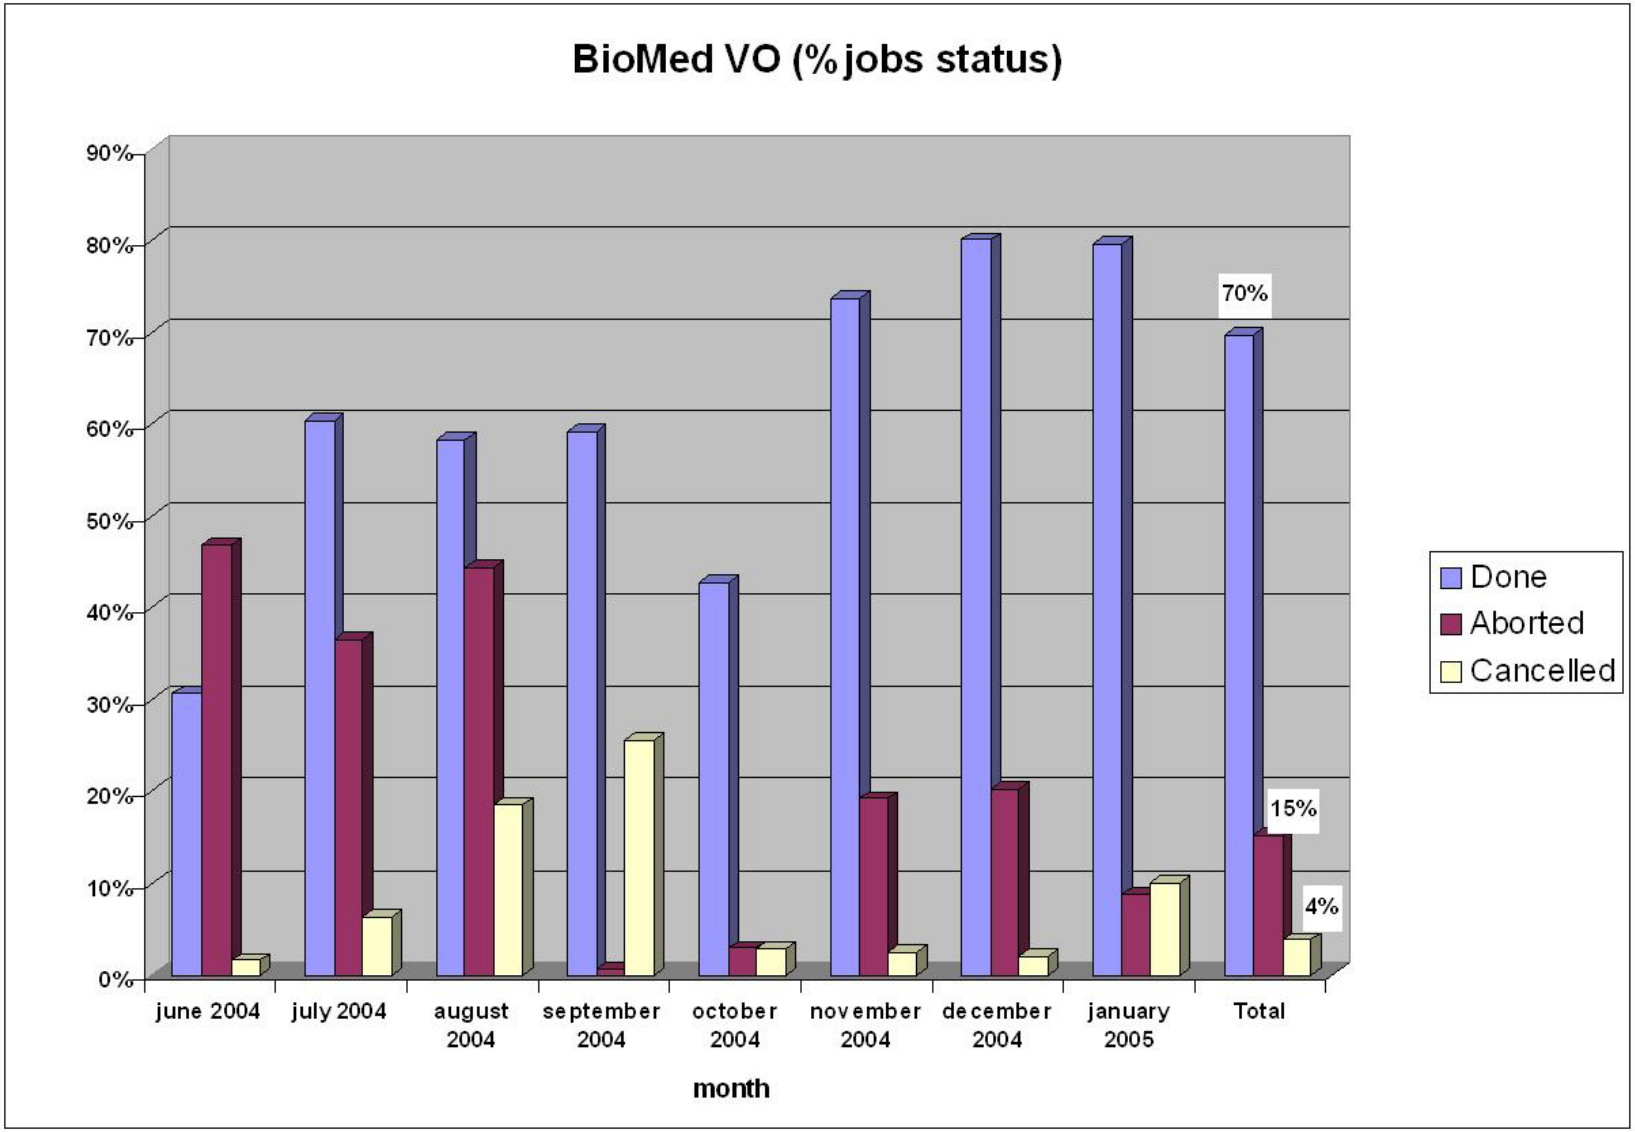

Biomedical applications

- 3 batch-oriented applications ported on LCG2
  - SiMRI3D: medical image simulation
  - xmipp_MLRefine: molecular structure analysis
  - GATE: radiotherapy planning
- 3 high throughput applications ported on LCG2
  - CDSS: clinical decision support system
  - GPS@: bioinformatics portal (multiple short jobs)
  - gPTM3D: radiology images analysis (interactivity)

  DEMO
- New applications to join in the near future
  - Especially in the field of drug discovery

Experience with LCG2 middleware

- **Two categories of applications had different levels of success**
  - Batch-oriented application (high performance): well adapted EGEE infrastructure, gridification has significant impact on performances
  - More dynamic applications (high throughput): gridification has been difficult with current middleware. Turn-arounds were needed to bypass some limitations
- **Still a high failure rate reported on LCG2 (order of 25%)**
  - Users tend to adapt manually their application (selection of sites to submit job, store data...)
  - Irregular through time (instability of the infrastructure)
  - This makes it difficult to estimate the failure ratio
- **The SA1-biomed interaction loop is being set up**
  - Significant improvement in feedback and solutions since Dec'04

Evolution of biomedical applications

- **Growing interest of the biomedical community**
  - Partners involved proposing new applications
  - New application proposals (in various health-related areas)
  - Enlargement of the biomedical community (drug discovery)

- **Growing scale of the applications**
  - Progressive migration from prototypes to pre-production services for some applications
  - Increase in scale (volume of data and number of CPU hours)

- **Towards pre-production**
  - Several initiatives to build user-friendly portals and interfaces to existing applications in order to open to an end-users community

Feedback to LCG-2 middleware developers and infrastructure

- **From HEP applications**
  - Experiment Integration Support group and Grid Applications Group produced documents summarizing problems encountered in use of LCG-2

- **From Biomed applications**
  - Very significant exchanges related to the set-up of the biomed VO and the deployment of relevant services
  - Request to use MPI

Identification and integration of new communities: EGEE virtuous cycle

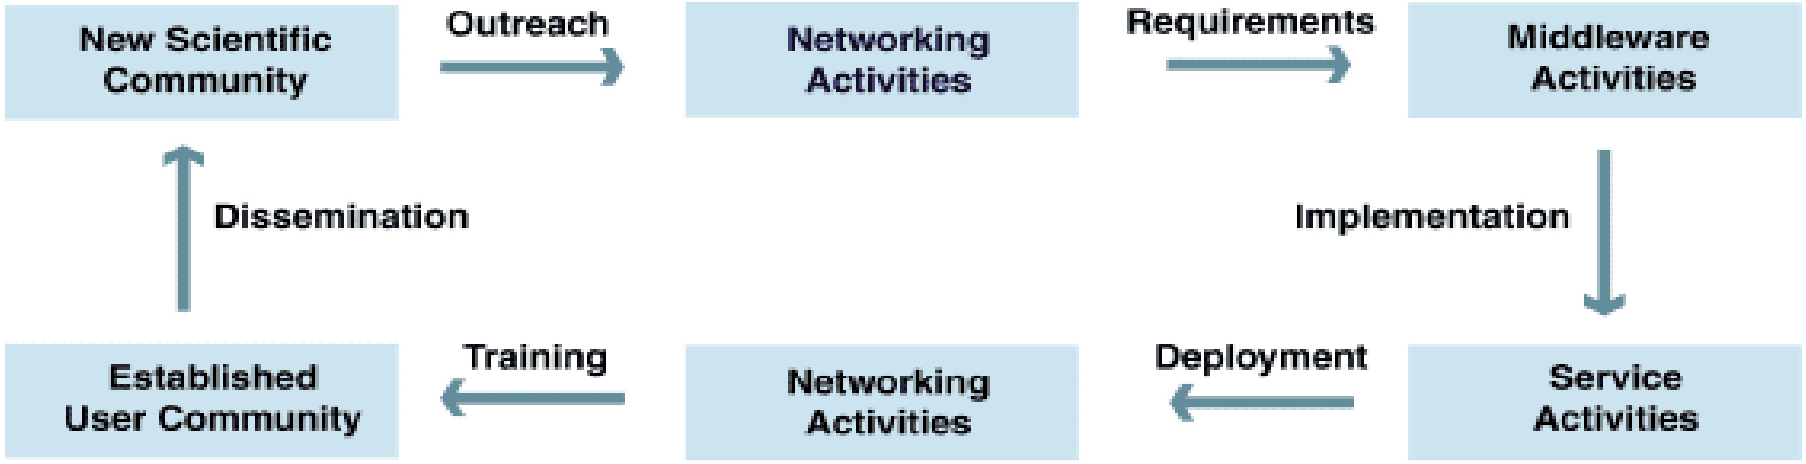

- Virtuous cycle concept is described in the project Technical Annex
- It describes the role of the different project activities to help new communities to successfully deploy applications on EGEE infrastructure
- As the first open multidisciplinary e-infrastructure in the world, EGEE has to invent the implementation of the virtuous cycle

New communities identification

- Through training, dissemination and outreach, communities already using advanced computing and keen to use EGEE infrastructure are identified
- These communities are encouraged to prepare a document describing their interest to use EGEE
- A scientific advisory panel (EGAAP) assesses and chooses among the interested communities the ones which seem the most mature to deploy their applications on EGEE

EGAAP

- **EGEE Generic Applications Advisory Panel is the entry door for new applications that want to be deployed on the EGEE infrastructure**

- **Important step in the EGEE virtuous cycle**
  - Encourages communities to submit a well documented proposal
  - Fosters discussion on the added value brought by the Grid to the applications
  - Points out needs and resources for migration and deployment for each application
  - Prioritizes the deployment of the selected applications
  - Monitors the progress of the selected portfolio

- **Participation in EGAAP of 5 external members is useful to reach out to new communities**

Summary of EGAAP activities

- **First call for proposals: limited distribution on May 17 2004**
- **First EGAAP meeting June 14 2004, at CERN**
  - 5 applications, 3 recommended for approval
    - Computational Chemistry
    - MAGIC, Astrophysics
    - Earth Science
- **EGAAP Recommendations approved by EGEE management on July 14**

- **Second call for proposals : widest distribution possible on Sep 28, 2004**
- **Second EGAAP meeting November 25, Den Haag**
  - 7 applications received, 4 recommended for approval
    - Earth sciences (Solid Earth Physics, Hydrology)
    - Cosmology (Planck collaboration)
    - Drug discovery (Molecular docking)
    - Search engines for digital knowledge (GRACE)
- **Recommendations approved by EGEE management on 17 December 2004**

EGEE Industry Forum

- **Objectives:**
  - To promote and disseminate Grid concepts towards industry and service groups
  - To raise the awareness of EGEE within industry
  - To encourage businesses to participate in the project

- **Members:** **interested companies having activities in Europe**

- **Activities:**
  - Organisation of a meeting twice a year
  - Quarterly newsletter
  - Participation to EGEE working groups (EGAAP, Project Technical Forum, EGEE Phase 2, Security group)
  - Internal Working groups
    - Technical aspects of Grid
    - Business models and economical aspects

Industry Forum members

MICHELIN
CEA
GRIDSYSTEMS
CREDIT LYONNAIS
CNES
BNP
PALLAS-INTEL
DASSAULT AVIATION
EADS CCR
SNECMA
STMicroelectronics Srl
Hewlett-Packard
CSTB
Daimler
Oracle
PSA
PLATFORM COMPUTING
GRIDXPERT
PECHINEY CRV
France Télécom

ACRI-ST
TOTAL
IFP
C-S
DATA SYNAPSE
THALES
LION Bioscience AG
BULL
Sanofi-Aventis
SOCIETE GENERALE
NICE
NOVARTIS PHARMA AG
EDF
ESI
Sun
ICATIS
CERFACS
MICROSOFT
British Telecom
CENAERO

Pôle Européen Plasturgie
SCAI
GENIAS Benelux
Arcelor
Fugitsu
IBM
HUTCHINSON
Gridwise Technologies
Compagnie Générale de Géophysique
SCHLUMBERGER
Datamat
AGENIUM Technologies
HLRS
Telefonica Spain
AIRBUS
DUTCH SPACE Paris Office
ORION LOGIC Ltd.
CSCS

From identification to integration

- Provision of training to the community application developers
- Identification of resources for new application deployment
- Assistance in interfacing applications to grid services
- Dissemination of information proactively addressing the needs (user support, middleware evolution,…)
- Monitoring of the integration process
- Provision of essential feedback to other activities - dissemination, middleware and management
- Definition of common application interfaces and tools

GILDA, an infrastructure for dissemination and demonstration

- **Goals**
  - Demonstration of grid operation for tutorials and outreach
  - Initial deployment of new applications for testing purposes

- **Key features**
  - Initiative of the INFN Grid Project using LCG-2 middleware
  - On request, anyone can quickly receive a grid certificate and a VO membership allowing them to use the infrastructure for 2 weeks
  - Certificate expires after two weeks but can be renewed
  - Use of friendly interface: Genius grid portal

- **Very important for the first steps of new user communities on to the grid infrastructure**

GILDA numbers

- 14 sites in 2 continents
- >1200 certificates issued, 10% renewed at least once
- >35 tutorials and demos performed in 10 months
- >25 jobs/day on the average
- Job success rate above 96%
- >320,000 hits on the web site from 10's of different countries
- >200 copies of the UI live CD distributed in the world

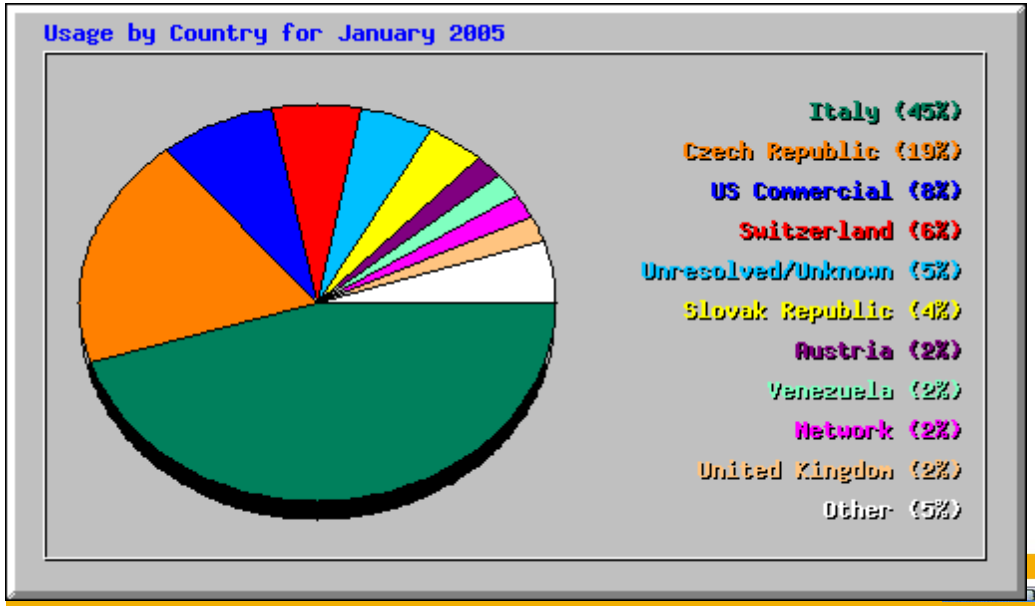

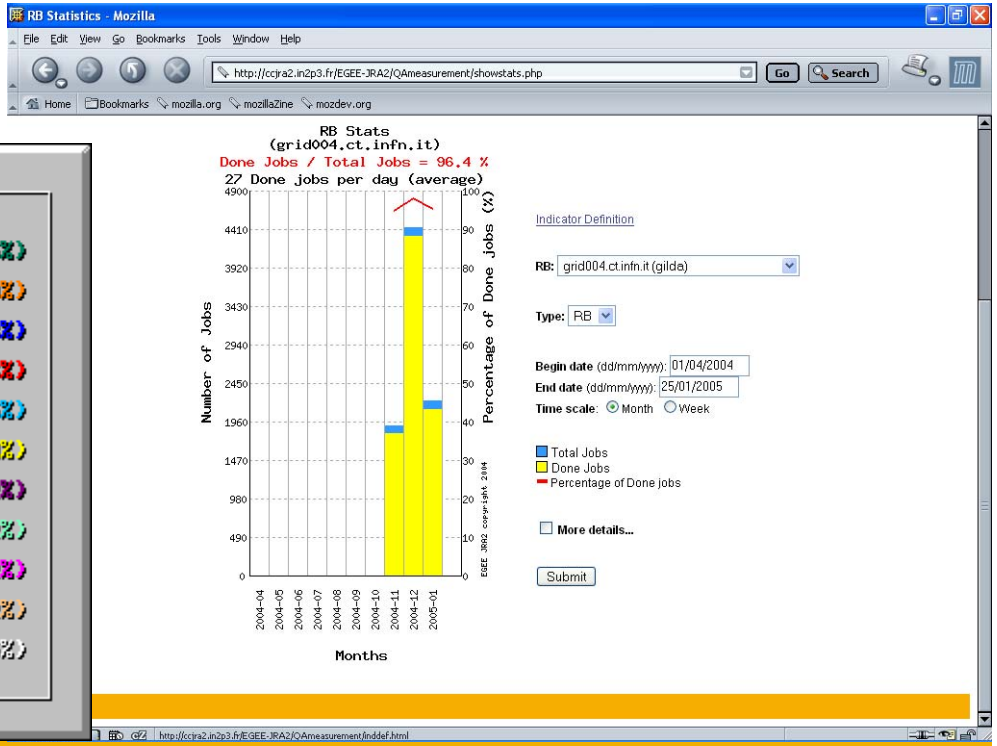

NA4 Applications and GILDA

- **7 Virtual Organizations supported:**
  - Biomed
  - Earth Science Academy (ESR)
  - Earth Science Industry (CGG)
  - Astroparticle Physics (MAGIC)
  - Computational Chemistry (GEMS)
  - Grid Search Engines (GRACE)
  - Astrophysics (PLANCK)

- **Development of complete interfaces with GENIUS for 3 Biomed Applications: GATE, hadronTherapy, and Friction/Arlecore**

- **Development of complete interfaces with GENIUS for 4 Generic Applications: EGEODE (CGG), MAGIC, GEMS, and CODESA-3D (ESR) (see demos!)**

- **Development of complete interfaces with GENIUS for 16 demonstrative applications available on the GILDA Grid Demonstrator (https://grid-demo.ct.infn.it)**

Transition from GILDA to Production Service

- **The transition from GILDA to the EGEE production infrastructure requires creating a new Virtual Organisation and deploying its services**

- **Each new Virtual Organisation requires the following services**
  - a VO administration service
  - a set of sites providing resources
  - potential access to a Resource Broker & Replica Location Service

- **The VO is administered by a VO manager who is a member of the associated scientific community. He is in charge of**
  - managing the list of VO users
  - monitoring the VO services and resources
  - informing the VO users of their availability

Earth Science Achievements & Issues

Achievements

- ESR (Earth Sciences Research) VO at SARA created in July 2004 and is functional using EGEE resources
  - 17 registered users from 6 countries
- The EGEODE (Expanding GEOsciences on DEmand) VO created at IN2P3 (Lyon) in mid-October for CGG and Geocluster partners
  - Preparation to migration to EGEE Production Service
- Important EGEODE application deployed on GILDA and demonstrated at the 2nd EGEE Conference in The Hague using the GENIUS portal

  DEMO
- Production of ozone profiles from the satellite experiment GOME and their validation by using LIDAR data run on EGEE production service

Issues

- Need secure access to data and metadata for authorised groups/sub-groups
- Access to licensed software

Number of jobs submitted by ESR VO members

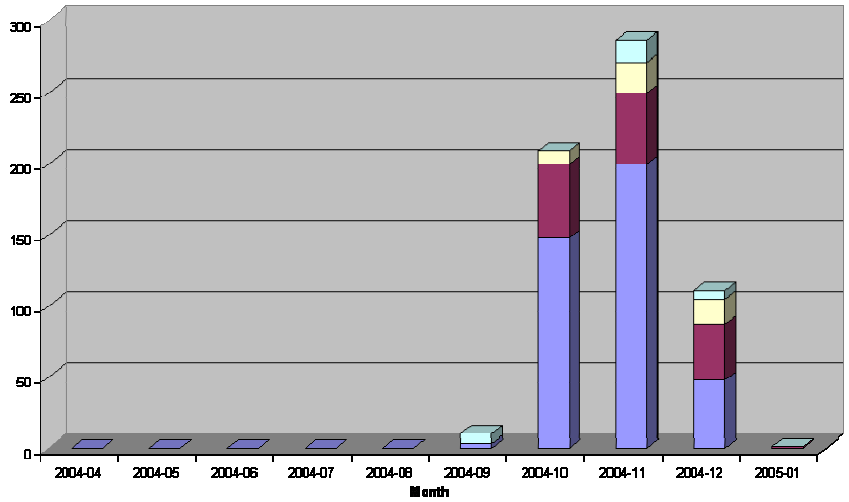

Computational Chemistry Achievements & Issues

Achievements

- Cluster of 13 nodes + CE + SE + VOMS server has been deployed in GILDA for dedicated use by CompChem.
- Grid based Molecular Simulator (GEMS) ported onto the GILDA test cluster and interfaced to GENIUS **DEMO**
- The CompChem VO has been activated
- Work in hand now to move to production service

Issues

- Requirements for interactive work
  - Outbound connectivity of worker nodes
  - Fast turnaround in jobs
- Access to licensed software

MAGIC Achievements & Issues

Achievements

- **A Magic Virtual Organisation already exists in EGEE**
  - VO server is hosted by SARA/NIKHEF
  - Successful first running in GILDA as well as in Crossgrid testbed using LCG-2 middleware
- **Developments underway for EGEE data challenge in early 2005**
  - CNAF will support the Magic VO with a Resource Broker
  - PIC will support the Magic VO with storage and the RLS
  - CNAF, PIC and GridKA will provide CPU
  - GILDA can be used for the first test as well

Issues

- **Education**
  - 'EGEE for dummies'
- **Getting extra EGEE resources for data challenge**
  - Precise 'process' definition and its execution

Major issues and their mitigation

- **Issue 1:**
  the planning and execution of the migration of applications currently deployed on LCG2, both HEP and non-HEP, to the new gLite middleware

- **Issue 2:**
  the provision of management and support structures for the integration of multiple user communities, and taking into account the significant increase in the number of EGEE active users

- **Issue 3:**
  the availability of security-enhanced services for data manipulation and job execution

Issue 1: Migration to gLite – the role of ARDA

- **High Energy Physics activity focused on the next generation applications**
  - Goal: allow physicists to do individual analyses of LHC data
  - Method: develop end-to-end prototypes based on the new gLite middleware
  - Resources: joint LCG-EGEE effort in ARDA (A Realisation of Distributed Analysis for LHC)

Migration to gLite: the role of ARDA

LHC Experiment	Main focus	Basic prototype component	Experiment analysis application framework	Middleware
LHCb	GUI to Grid	GANGA	DaVinci	gLite Lightweight Middleware for Grid Computing
ALICE	Interactive analysis	PROOF ROOT	AliROOT	
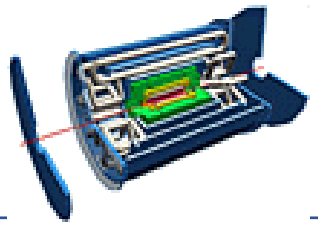	High level services	DIAL	Athena	
CMS	Exploit native gLite functionality	Aligned with APROM strategy	ORCA	

Migration to gLite: the role of the NA4 Test team

- **Goal: develop and execute test cases corresponding to application use cases**
- **Strategy:**
  - Propose test cases based on use cases collected from application developers (http://marianne.in2p3.fr/egee/testgroup/testcase/Tables/)
  - Design and implement a test suite based on the test cases and compliant with EGEE test strategy
  - Perform tests on the pre-production service and the production service
- **Milestone MNA4.1 accepted**
- **Status**
  - Implementation of a new set of components: improved modularity, easier test building, better integration in a test framework, new functionality
  - Participation to EGEE common testing activities

Issue 2: Enabling the virtuous cycle

- **Provision of management and support structures for the integration of multiple user communities is a project wide challenge**
- **It requires information flow between project activities**
  - SA1: infrastructure
  - NA2: dissemination
  - NA3: training
  - NA4: application deployment and support
- **Urgent needs identified:**
  - Strong interface between VO managers and infrastructure operation management
  - User support structure needs to handle a growing number of users
  - Single entry point to the project for external user communities

Issue 2: creation of inter-activity groups

- **Major role of the NA4/SA1 joint group**
  - already described in Ian Bird's talk
  - Understands the detailed requirements for new communities joining EGEE
  - Assists in negotiations for resources for the new applications community
  - Ensures that the full services of the infrastructure will be available to the new VO
  - Ensure that the new VO demonstrates an appropriate commitment to the project
  - Negotiate that sites provide a slice of their resources for test purposes in order to encourage applications to join EGEE
  - Also broker for existing applications needing more resources
- **Reflection started on the design of the user support structure (SA1)**
- **Reflection started on the design of a single entry point for EGEE external users (NA2/NA3/NA4/SA1)**

Issue 3: Security-enhanced services

- **Security-enhanced services for data manipulation and job execution are critical for industrial partners and non-HEP scientific applications**

- **gLite 1.0 will provide new security-enhanced services**
  - VOMS for VO management
  - The File and Replica Catalog provides support to ACLs
  - These services must be tested by the biomedical pilot applications

- **Additional support needed**
  - ACL support at Storage Element level

- **Secure access to licensed software – will be investigated via specific cases with members of the Industrial Forum**

Plan for next period

MNA4.2	M12	First external review of Applications Identification and Support with feedback
DNA4.3.2	M15	First revision of EGEE Application Migration Progress report
DNA4.3.3	M21	Second revision of EGEE Application Migration Progress report
MNA4.3	M24	Second external review of Applications Identification and Support with feedback
DNA4.4	M24	Final Report of Application Identification and Support Activity

- No change w.r.t. to TA
- Lack of resources to address issues identified in relation to the integration of new user communities

Summary

Top accomplishments

- The successful deployment of several biomedical applications
- The successful outreach to new generic communities through a well established process, providing education and application migration for the new application areas selected by EGAAP, using GILDA and GENIUS as tools
- The High Energy Physics community has made an extensive use of the production infrastructure
- The demonstration of prototype analysis systems using gLite for all 4 LHC experiments

Major challenges for the coming months

- Planning and execution of the migration to gLite of applications currently deployed on LCG2
- Provision of management and support structures for the integration of multiple user communities taking into account the significant increase of the number of EGEE active users
- Availability of security-enhanced services for data manipulation and job execution