



Enabling Grids for  
E-science in Europe

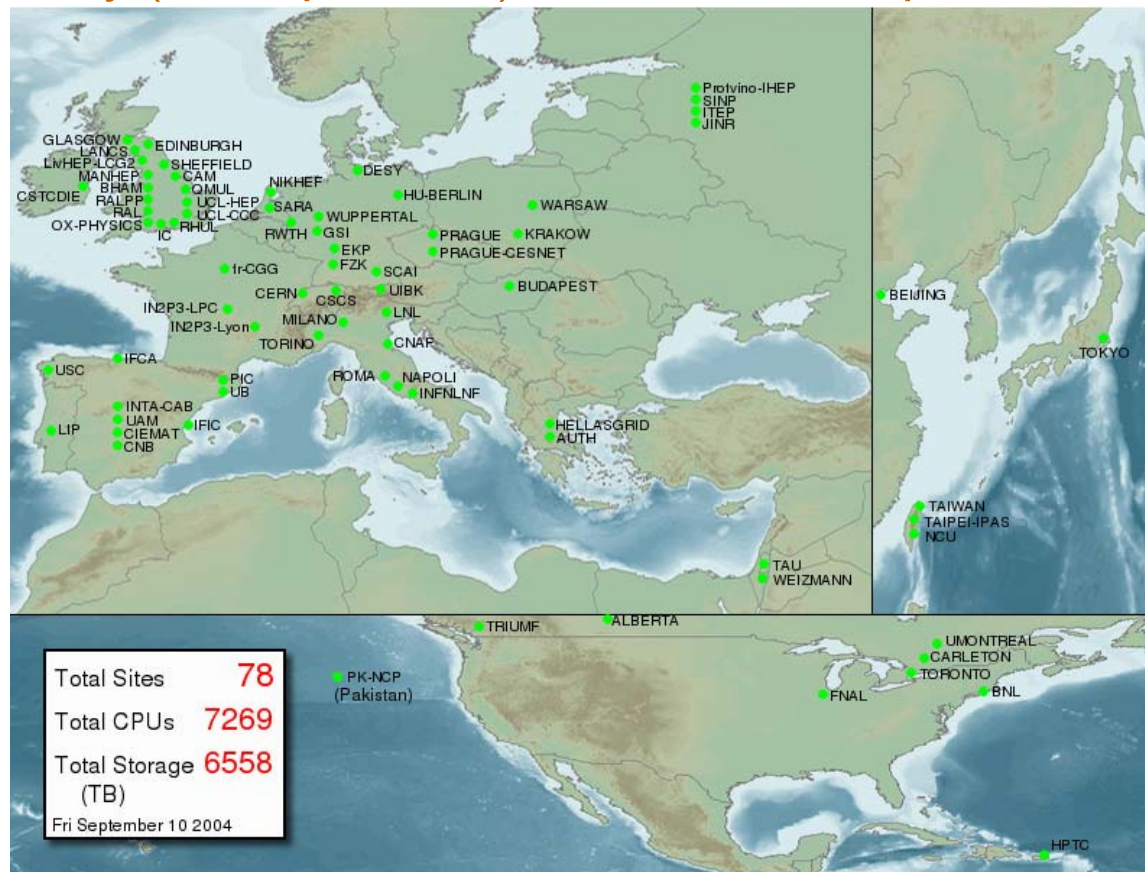
# SA1 All Activity Meeting 13 September 2004

Ian Bird, Cristina Vistoli  
and input from ROC managers



# Accomplishments since last AAM

- Production service is fully operational and continues to grow
  - Currently (10 September) 78 sites, 7269 cpu



# Accomplishments since last AAM

- ROCs are set up and starting to take up their support responsibilities:
  - This must become fully developed during the rest of the year
  - Effort started in developing common support infrastructures
- CICs
  - Nick Thackray is coordinating – developing planning for CICs to evolve out of LCG GOC
  - Lyon CIC is supporting bio-med
- LHC experiment data challenges have been running for past 6 months
  - Required very high operational maintenance load

- Completed Execution Plan (DSA1.1)
  - Of course with regional contributions and help
- Completed 1<sup>st</sup> quarterly report
- Most (huge!) effort has been fully devoted to running operations during the ongoing LHC data challenges
  - This load must be picked up by the ROCs and CICs now – this will be a key test for EGEE operations
    - Issue – How can we provide 24x7 operations with staff hired for 8x5 (true for all federations)
  - This load will remain and we will start to see new VOs also
  - Most new sites have so far been supported directly from CERN
    - ROCs must pick up this load
- Operational security group – led by Ian Neilson
  - Planning better incident response and daily security procedures
- VO management
  - Provided assistance for new VOs, registration procedures, etc
- Provided staff to help with training courses
- Staffing is complete
  - 1 new project associate to support SEE-grid started August

- **ROC**

- The GridPP Grid is now 14 sites (end August) running LCG2 software.
- The National Grid Service sites have not yet been certified.
  - They are running VDT, the GLUE schema and have a BDII and tests with RB
  - Need to get all running on RHEL
- GridIreland running LCG-2\_2\_0 internally on testbed including national servers (RB/BDII/MyProxy/RLS)
  - Set to deploy LCG\_2\_2\_0 on Grid-Ireland infrastructure this week (Week start 6-Sept-04)
  - Ongoing porting effort to IRIX and AIX in Ireland(unfunded by EGEE)
- Staff from the UK helpdesk have been part of the SA1 work with GGUS at Karlsruhe.
- **Security** The LCG Security Group → Joint Security Group (JSG); JSG is led by UK/I ROC.
  - Site security requirements were collected and fed into the JRA3 requirements analysis.
- **Training** – A course on installing LCG2 for sysadmins was run at Oxford in July. This will be repeated later in the year. LCG-2 dissemination, planned 3-day workshop in Ireland September 20th-22nd 2004

- **CIC**

- **Monitoring**
  - Development of regional monitoring. Tailored maps for each of the ROC showing .Extend with other info (inc accounting)
  - Replaced manual daily reports (sent to LCG Rollout) by an automated process involving RSS feeds
- **Accounting**
  - LCG Accounting package has been sent to C&T team and we are in email discussions with them to resolve any problems/issues.

- **Staffing**

- The UK is almost up to strength but there will still be recruitment of some extra unfunded effort.
- Ireland operating at full staff levels since 26th July 2004.

- **ROC**
  - 18 sites (~ 1200 CPU's) running INFN-GRID (LCG2 based) middleware participate to the INFN-GRID infrastructure
  - The National Grid Service
    - INFN-GRID Release:
      - 15 VO's supported (alice, atlas, babar, bio, cms, cdf, enea, gridit, inaf, infngrid, ingv, lhcb, theophys, virgo, zeus)
      - Full VOMS features
      - Resource usage metering and job monitoring available
    - 3 Resource Brokers/BDII with visibility of the all the national grid resources
    - 1 GridIce server for monitoring services, resources and jobs
    - 1 Replica Location Server for non-LHC VO's
    - VOMS and VO-LDAP server for national VO's
    - Certification service: test zone for site certification
  - Plan to add Resource Centers of italian EGEE partners
  - Operation and Support
    - Web-based tool to manage downtime advices and inform users and site managers about scheduled maintenance activities for sites and services
    - Scripts for 'on demand' site certification to verify site behaviour
    - Remote control for partially unattended sites
    - Ticketing system: 495 Ticket in 6 months time, mainly used during upgrade phase to exchange info between site managers and ROC
- **CIC**
  - EGEE-wide Grid services
    - EGEE Resource Broker and BDII for BIOMED and used by LHC-B-DC
    - ATLAS VO-specific Resource Broker and BDII with sites list managed by CIC people and controlled by experiment people
    - VOMS and RLS server for EGEE generic application VO
  - Monitoring/Accounting
    - GridICE development and integration in the LCG release
    - Job monitoring and usage metering graph per VO and site
- **Staffing:** almost complete

- Resource Centres:
  - Lyon (full function, not all machines yet).
  - Clermont (full function, not all machines yet).
  - CGG (Compagnie Générale de Géophysique) operational.
    - Will probably request addition of VO ESI (Earth Science Research for the Industry) at some time.
  - LAL Orsay coming up, partially functional (UI, BDII).
  - IPSL site coming up at Paris (Earth Science Research).
    - Requesting addition of VO ESR. VO server is at the moment situated at SARA (NL).
- ROC:
  - Assisted in installing and configuring the majority of the RCs mentioned.
  - Participates in GGUS task force
  - Provides a contact for NA3 in France for future user training events
- Biomed VO basically operational:
  - VO server at Lyon. RB at CNAF, 2nd RB in setup at Lyon, 3rd RB in setup at Barcelone
  - RLS in setup at Lyon
  - additional sites in Israël and Spain (Madrid).
- CIC activity:
  - VO integration: draft paper available. Is part of larger subject:
  - VO administration, under work on CIC level.
- Staff situation:
  - The number of FTEs was erroneously 26 in the last AAM report, should have been 24, is now 25.
  - About 50 people involved to some degree in SA1. The additional FTE is due to the recruitment of one person for the CIC activity in Lyon.

- **Status**
  - Successfully started to set up the organizational and operational infrastructure.
  - The operational issues tackled in this period include:
    - supporting the 3 production clusters (2 in IL and 1 in GR) and
    - preparing the clusters in BG, CY, GR, IL and RO;
    - setting up pre-production and test clusters;
    - Defining technical solutions for helpdesk and monitoring,
    - running the SEE portal;
    - running the CAs, etc.
  - Most of these were preparation activities. In the period from now until the end of the year SEE should see production sites supporting various EGEE VOs. SEE should have more production and pre-production clusters joining the infrastructure, as specified in the Execution Plan, and prototype monitoring and helpdesk solutions will be under development.
- **Issues**
  - The partners in the region find that the reporting procedures, such as quarterly reports, timesheets, etc, have proven to be very heavy-handed and time-consuming.
  - We hope that these procedures will be relaxed
- **Staffing:**
  - People Total (F+U): 45
  - People Total (F): 33
  - FTEs total (F+U): 15.68
  - FTEs total (F): 12.49



# South West

- 9 sites connected



- Relevant contribution to LCG Data Challenges from several sites using the Grid (CMS, Atlas and LHCb). See GDB transp.
- SA1 Partner changes: IFAE becomes PIC as a Joint Research Unit, and one new partner join: Telefonica I+D. The execution for SA1 plan has to be modified to include Telefonica I+D.
- Numbers: 6 SA1 partners, 9 resource centers, 34 people involved (~20 FTE).

- Resource Centres
  - 9 sites running with 237 WNs
  - Several partners wait for new hardware – larger amount of CPUs will be available soon
  - Various EGEE VOs supported
  - Most of partners are experienced in LCG
- ROC
  - Organization according to Execution Plan is established
  - Two sites in the pre-production service
  - Test procedures for RCs in progress
  - Monitoring and helpdesk is being organised
  - The national CAs are running
- Staffing:
  - 48 persons engaged (24 funded and 18 unfunded) mostly part-time
  - Staffing has not reached maximum level yet. Additional recruitments are planned

# Northern Europe



- SARA (Netherlands Belgium)
  - Resource centers NIKHEF and SARA have both upgraded their EGEE resources to LCG2\_2\_0.
  - Several new VOs have been added to the local LDAP server: *esr* with subgroups, *astrop*, *magic*. Also RLS services have been setup for these VOs.
  - Scripts have been developed to facilitate the introduction of a new VO to the local configuration. These scripts are available for use in EGEE and will be included in future software releases.
  - The compute resources for EGEE at SARA use TORQUE with the MAUI scheduler since August 3, NIKHEF already uses this batch environment for their resources.
  - And loads of sysadmin work has been done these months to keep jobs flowing.
  - Belgium Resource Centers are installing LCG2 on their clusters and they will build a local infrastructure with basic services.
  - New equipment is expected to be added in the coming months. Integration of resources of NLGrid and BEGrid will be the next step.
- SNIC (Nordic, Estonia)
  - Substantial amount of the ATLAS Data Challenge workload has been taken on the three Swedish resource centers (SWEGRID).
  - Cluster has been set-up ready for installation EGEE pre-production environment.
  - LCG2\_2\_0 deployment is underway at three resource centers (Umeå, Linköping and Stockholm).
  - LCG2\_2\_0 installation adapted for Debian is underway.
  - A website for support and internal project use has been set up for the NE ROC <[www.egee-ne](http://www.egee-ne)>.
- Staffing:
  - SARA: 1.25 FTE lower than expected in TA
  - SNIC: staffing is complete


# Germany + Switzerland

- **ROC**
  - **Distributed ROC** with partners DESY, FhG (2 Institutes), FZK, and GSI
    - all 5 sites running LCG2 software (~1000 CPUs in total); 3 additional non-EGEE sites registered; upgrade to LCG2\_2\_0 started
    - RB, RLS, VO-server, CA running at different sites
    - support for 13 VOs: Atlas, Alice, CMS, LHCb, BaBar, CDF, Dzero, DESY, H1, HERAB, HERMES, LC, ZEUS
  - **Rotating mw support** to gain experience: at FZK until end of August; now at FhG/SCAI
  - **User support**
    - Global Grid User Support running at GridKa
    - a support task force was set up (members of all federations; led by FZK). Goal: propose methods and workflows for EGEE user support, including already existing regional/national support teams and tools; several meetings by phone, meeting in person at FZK in August; draft concept submitted to ROC managers in July, final concept due by end of September
  - **Training & Dissemination**
    - preparing open day for the general public at FZK (Sept 18) with public talks on Grid, guided tours through computing centre, performances etc.
    - GridKa School '04 (Sept 20-23) together with NA3 with user tutorial, admin installation course, tutorials for HEP applications and core software
- **Staffing**
  - Almost complete, but already lost one person; have to fill this gap during next weeks
- **Issues**
  - Requests from swiss users for mw support, but there is (yet?) no official swiss contract partner in EGEE to contribute

# Russia



- **ROC:**

- Four of eight sites were connected. 
- LCG2 based middleware participate to the Russian Data Intensive GRID (RDIG) infrastructure
- The National Grid Service
  - RDIG Release:
    - 5 VO's supported (alice, atlas, cms, lhcb, dteam)
    - Resource usage metering and job monitoring available
  - 1 Resource Brokers/BDII with visibility of the all the national grid resources
  - Certification service: RU-zone for site certification
- Operation and Support
  - Web-based tool to manage downtime advices and inform users and site managers about scheduled maintenance activities for sites and services
  - Remote control for partially unattended sites
- Three (PNPI, RRC KI, IMPB) sites have already installed the LCG-2\_2\_0;
  - Issue – VO Biomed is not supported in Russian CIC
- The other one site (KIAM) will installed the LCG-2 at Month 15.
- ROC is set up now and helps Russian RCs;
  - Issue – 24x7 operations

- **CIC:**

- EGEE-wide Grid services
  - EGEE Resource Broker for Russian RC
  - BDII for Russia
  - Certificate Authority for Russian users/sites.

- **Staffing:** almost complete

- **Nearest plans:**

- **Add new RC: PNPI, RRC KI, IMPB and other**
- VOMS and VO-LDAP server for national VO's
- Monitoring/Accounting
  - Gridlce server for monitoring services, resources and jobs
- Ticketing system to exchange info between site managers and ROC

# Status on PM6 Milestones & Deliverables

- MSA1.1
  - “Initial Pilot Grid operational with 10 sites. The ROCs and CICs in place.”
  - This milestone is easily achieved with the existing infrastructure and service. The ROCs and CICs are organised and in place.
  
- DSA1.2: “Release notes to accompany MSA1.1”
  - Format:
    - The release notes will be a short covering document (~5 pages) giving the overall framework and pointing into the existing web-based documentation.
  
  - Scope and TOC:

# DSA1.2 TOC

- The EGEE grid infrastructure
  - The national grids and LCG
  - The regions/federations
  - The production service and the test zones: national grids and VO grid services
  - FAQs
- Middleware overview and contents
  - Overview - what it does, does not do - what it contains: VDT, EGEE, LCG stuff etc
  - Users' Guide
  - Scenarios/example
  - Release notes –
    - Known problems (e.g. docs that Flavia etc are currently writing), Changes since last release
  - FAQs
- Installers
  - Installation guide
  - Notes on installation methods - lcfg, manual etc
  - FAQs
- Operations support and management guides
  - Communications
  - Support
  - Grid management
  - known problems and limitations
  - FAQs
- Security
  - users and system managers guides
  - Procedures etc
- Documentation overview and index

# Top 5 Risks

- gLite middleware is too late and/or too complex
  - this I think is a real risk - we saw exactly the same one year ago. The complexity can cause other problems - it makes it hard to use and this discourages users and so the project does not become widely used except for HEP. This would be a disaster.
- security incidents break existing trust between sites and cause sites to withdraw,
- failure to attract new applications and user communities
- failure to integrate existing national and regional infrastructures (Nordugrid!).
- failure to work with other international grid projects



# Issues related to other activities

- NA3:
  - Already 2 proposed SA1-NA3 liaison people withdrew, falls back to CERN
  - ...
- NA4:
  - SA1/NA4 group being set up
  - Issue: need to see organised teams in these VOs with good plans and goals who want to make their apps grid-enabled.
- JRA1:
  - Continuous ongoing interactions at several levels
  - Concern about gLite delivery timescale
  - Concern about manageability and operability (it *must* be better than LCG-2) otherwise to operations load will be unsupportable
  - Bottom line: we need gLite in our hands (on pre-production)
- Project:
  - Time consuming and heavy reporting (timesheets, quarterly reports,...)
    - Brought up many times by all federations

# Priorities until den Haag

- ROCs really take over front line support
- CIC's take the operational support load
  - We need well-defined operations managers in each who take the responsibility to be “the operations manager” for a week at a time, rotating through the CICs
- Accounting and more complete system monitors
  - This is essential (M9 deliverable) and needed by CICs
- Pre-production service in place
  - With first gLite components (WMS, CE)

# Pre-production service

- Nick's talk

# Migration to gLite

- Pre-production service will start with LCG-2
  - Promised by end of September to have gLite WMS and CE
    - This will be deployed also on pre-prod
- Expect LCG-2 and gLite to co-exist – not to interoperate?
  - But does this mean gLite does not interoperate with other grids? (step backwards)
- Impossible to plan migration until we know more about the actual implementations
  - And have tried them together with LCG-2
  - Need to understand deployment issues – dependencies, compatibilities, etc.
  - This must be exposed to users, deployers, admins, security people
- Hard to imagine that we can have enough understanding to build a migration plan by Christmas

# Support for non-HEP apps

- Many non-HEP (and non-LCG HEP) VOs are supported regionally already
  - Several with interest to get more resources
- Biomed is first non-HEP app to be deployed EGEE-wide
  - We need to see strongly organised groups within new VOs that we can work with directly to get them going:
  - New VOs must create teams dedicated to doing this, and who feel it is their responsibility to make it work; this must happen soon
  - We would like to see some realistic goals defined by the new VO groups as to what they want to achieve
- SA1/NA4 group can be a vehicle for this (ref – to paper)
  - 1<sup>st</sup> meeting this week
- Integration of new VO's
  - Draft paper from France based on experience with setting up biomed <https://edms.cern.ch/document/488885>
- Most RCs want to support the VOs they are funded to support
  - Resources not being made available to the other VOs.
  - We would probably like sites to make their resources available to all support EGEE VOs
  - Need to start mandating “10%” queues for such use (for example)
  - SA1/NA4 group should agree this policy

- Providing “24x7” support with “8x5” staff → How?
- We need to see strongly organised groups within new VOs that we can work with directly to get them going:
  - New VOs must create teams dedicated to doing this, and who feel it is their responsibility to make it work; this must happen soon
- Reporting/documentation load is high – and directly competing for staff who are fully utilised running operations
  - ROCs and CICs must rapidly take the load from CERN and GOC teams
  - Can we stop operations for the review?
- gLite – SA1 has seen no components yet (for good reasons)
  - But already this means that fully functional system at M14 is in doubt – we know it takes 6 months to make things “production” quality