



Enabling Grids for E-science

SA1 Report

Ian Bird
CERN

EGEE-II
All Activity Meeting

24-25th August 2006

www.eu-egee.org



- **Progress and status**
- **Execution plan – WBS status**
- **Issues**
- **Plans**
- **Progress on recommendations from EU reviews 2 & 3**
- **Deliverables/milestones & Summary of key deliverables**
 - DSA1.1 GGUS plan
 - DSA1.2 OAG procedures and policy
- **Report on OAG**
- **Internal communication**
- **Summary**

- **Since April/May**
 - Highlights
 - Deployment of gLite-3.0
 - LCG Service Challenge 4 throughput and service phases
 - Operations workshop – focus on interoperability with OSG
 - (non grid) security incident
 - Main WBS tasks:
 - Production service – versions, deployment progress
 - *Ops workshop 19-20 June*
 - *Support for Service Challenges*
 - PPS – versions, problems, sites
 - *gLite-3.0 etc testing; response for upgrades much improved*
 - Deployment and support
 - Grid Operator on Duty progress
 - Security – incident response, vulnerability, JSPG
 - Support
 - *GGUS, VO support*
 - Interoperation
 - *Many activities*
 - SLA – site policy
 - OAG
 - *Resource access policy*
 - Input to TCG

19-20th June at CERN

- **Theme: Readiness for LHC Commissioning. (Include OSG-EGEE interoperation):**
- <http://agenda.cern.ch/fullAgenda.php?ida=a062031>
- **Solve Current problems; Increase involvement; Identify lacks.**
- **Format:**
 - Introductory plenary followed by parallel sessions (2 chair persons per session: one from OSG and another from EGEE)
 - Release and deployment process
 - User support
 - Operations process
 - Site/service validation
- **Closing plenaries:**
 - Vulnerability group
 - OSG risk assesment
 - Summary of workshop
- **Attendance: 80 people**

Follow Up:

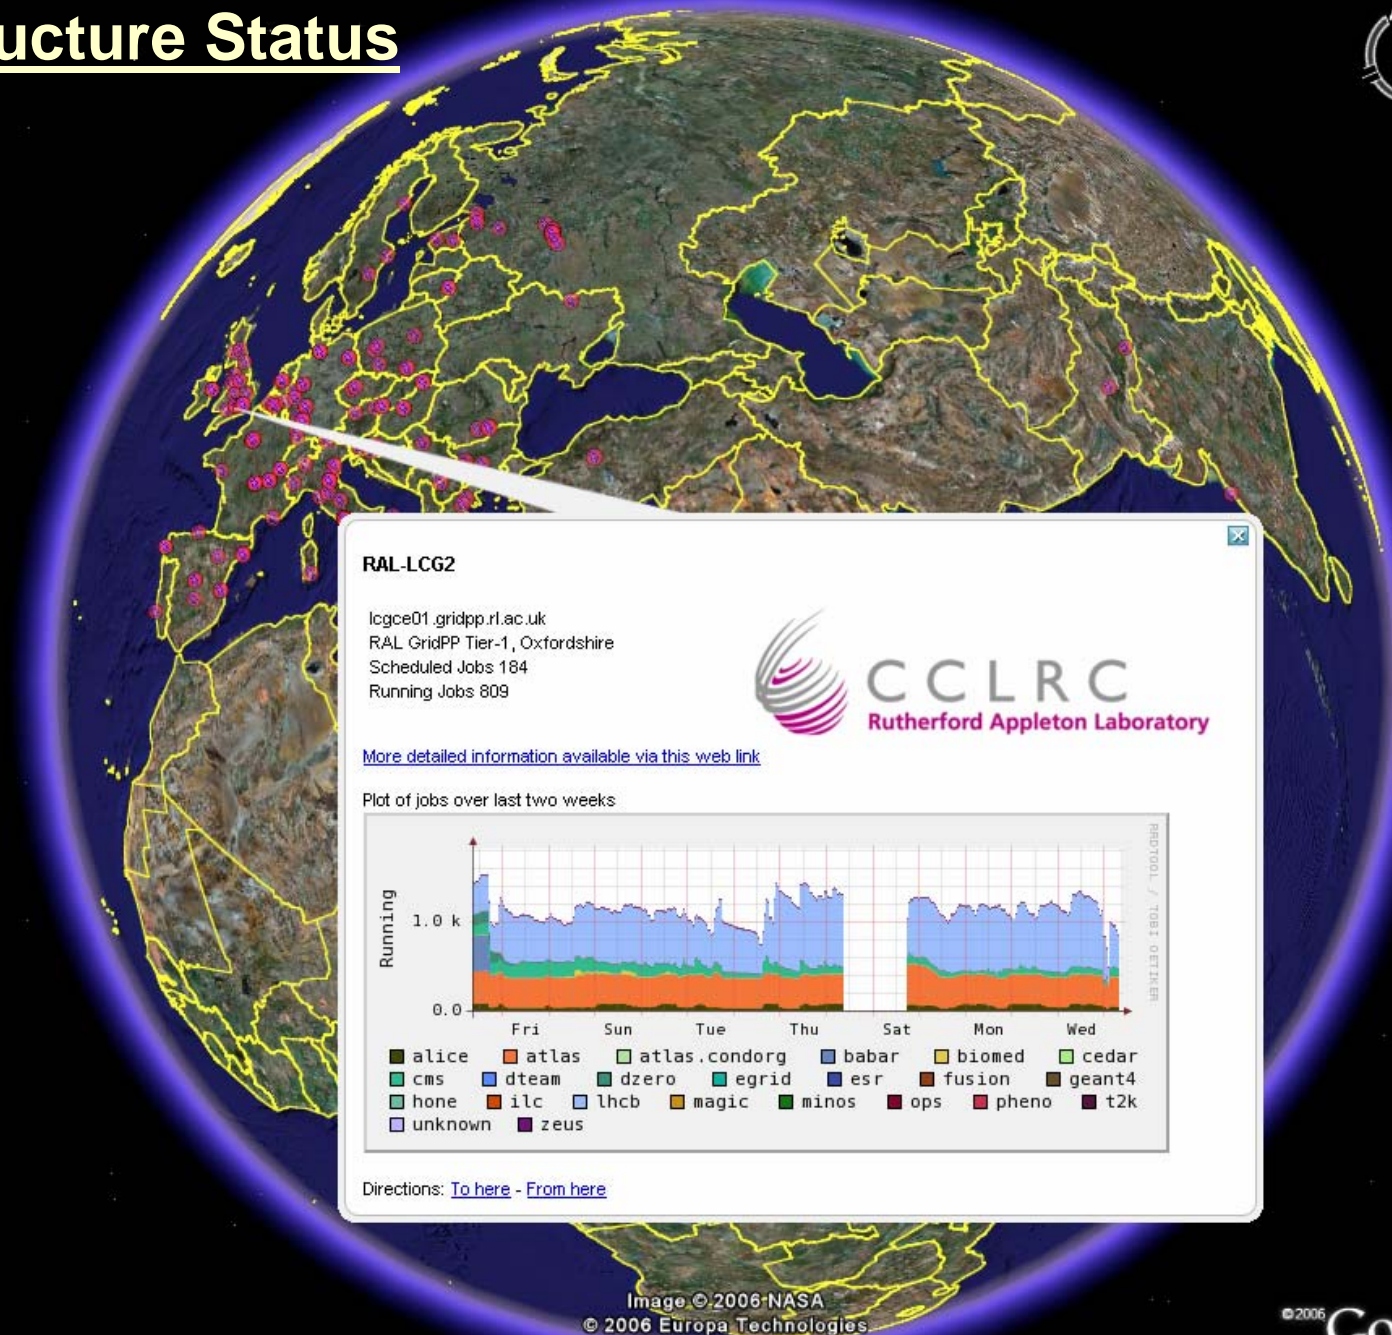
- **Action items being followed up in operations meetings and ROC managers meetings**
 - Already resolutions to several of them
- **Agreed by all that these workshops are extremely useful –**
 - especially for solving interoperability/interoperation issues
- **Next meeting –**
 - next “Spring” (March) in Indiana?

- **gLite-3.0 available for deployment on May 4**
 - 3 days later than scheduled
- **Deployed to large sites during May (all LCG Tier 1 sites)**
 - General deployment followed, some installation issues were resolved
 - Updates 3.0.1 and 3.0.2
 - Despite the significant change in the release the production system was upgraded and continued to provide production service
 - 3.0.1 22 sites
 - 3.0.0 112 sites
 - 2.7.0 40 sites
 - 2.6.0 1 site
 - ... 17 sites either down or very old version
- **Still do not have production quality gLite-flavour WMS/CE**
 - Concentrated effort with CMS/ATLAS to work on RB – 1 month and still ongoing
 - Will need to do similar effort with the CE
 - → this is testing in production!!!

- **A security incident occurred that affected many (all?) of the large HEP sites and several others**
- **This was not grid-related (i.e. did not involve grid credentials or services)**
- **It was a serious incident and had been ongoing for several weeks when it was discovered**

- **But: it did test the response and coordination procedures successfully**

Infrastructure Status



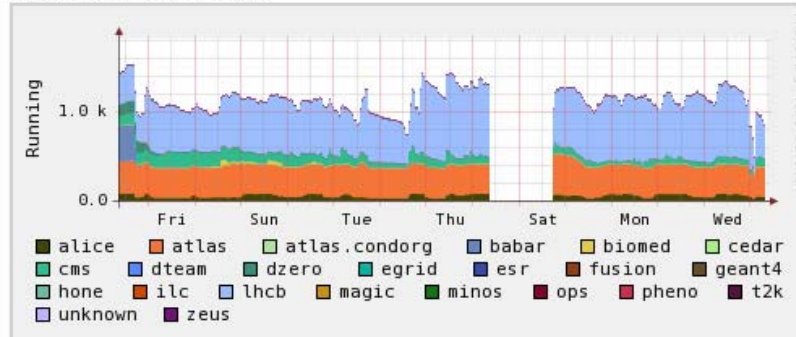
RAL-LCG2

lcgce01.gridpp.rl.ac.uk
RAL GridPP Tier-1, Oxfordshire
Scheduled Jobs 184
Running Jobs 809

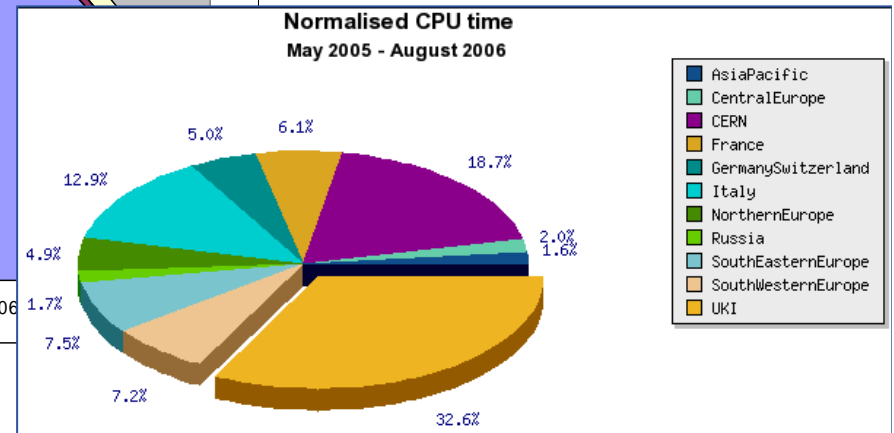
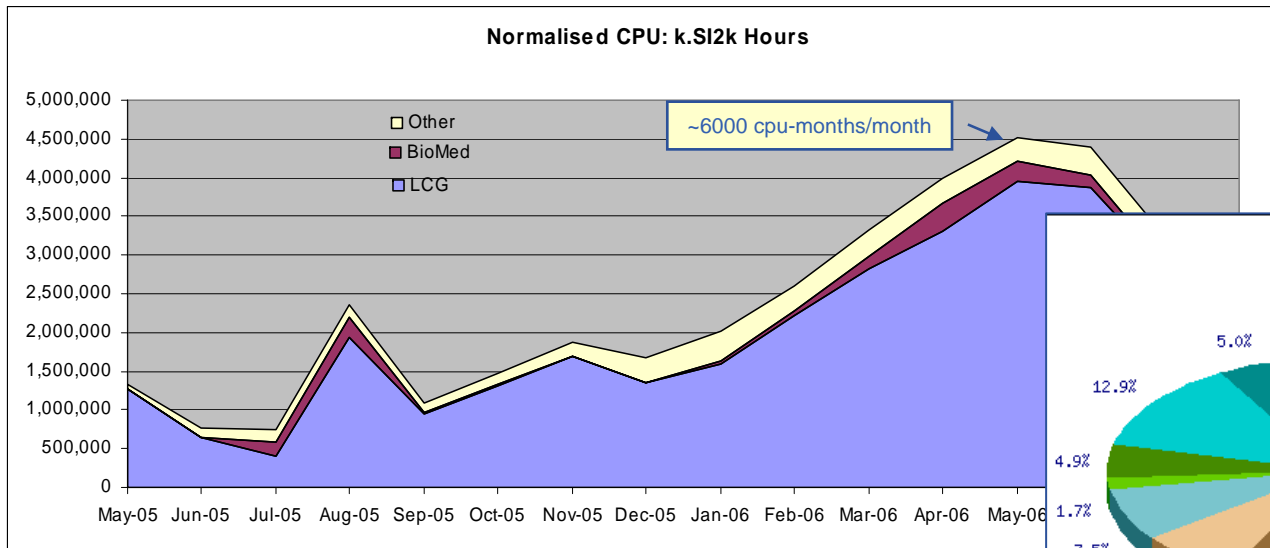
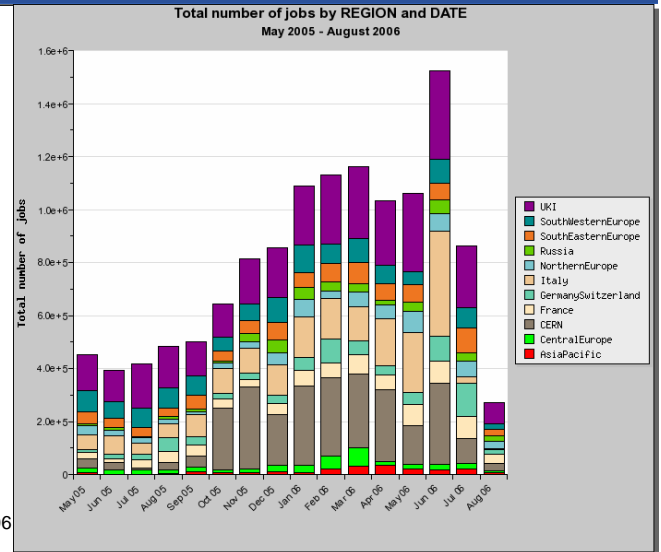



[More detailed information available via this web link](#)

Plot of jobs over last two weeks



Directions: [To here](#) - [From here](#)



<i>Region</i>	<i>#countries</i>	<i>#sites</i>	<i>#cpu</i>	<i>#cpu DoW</i>	<i>disk (TB)</i>
<i>CERN</i>	0	1	4400	1800	770*
<i>UK/I</i>	2	23	4306	2010	310
<i>Italy</i>	1	27	2800	2280	373
<i>France</i>	1	10	2316	1252	300*
<i>De/CH</i>	2	13	2895	1852	280*
<i>Northern Europe</i>	6	16	2379	1860	64
<i>SW Europe</i>	2	13	956	898	16*
<i>SE Europe</i>	8	26	1101	1885	30
<i>Central Europe</i>	7	21	1584	1163	70
<i>Russia</i>	1	15	515	445	38
<i>Asia-Pacific</i>	8	19	840	751	72
<i>North America</i>	2	8	4069	[4069]	229
Totals	40	192	28161	20265	2552

* Estimates taken from reporting as IS publishes total MSS space

TSA1.1: Operate a production and pre-production service

TSA1.1.1: ROC management (258)

TSA1.1.2: Pre-production service site (290)

TSA1.2: Middleware deployment and support

TSA1.2.1: Coordination and support for middleware deployment (516)

TSA1.2.2: Regional certification of middleware releases (100)

TSA1.3: Grid Operations and support

TSA1.3.1: 1st line support for operational problems in region (408)

TSA1.3.2: Oversight and management of operational problems (294)

TSA1.3.3: Run essential regional grid services (270)

TSA1.3.4: Weekly operator on duty support (240)

TSA1.3.5: Grid services for infrastructure or VOs (221)

TSA1.4: Grid security and incident response

TSA1.4.1: Grid incident response coord in region (150)

TSA1.4.2: Security vulnerability and risk analysis (68)

TSA1.4.3: CA management (99)

TSA1.4.4: Coordinate JSPG (18)

TSA1.4.5: Coordinate EUGridPMA (20)

TSA1.5: VO, application, and user support

TSA1.5.1: GGUS (60)

TSA1.5.2: Call centre, helpdesk for ROC (330)

TSA1.5.3: VO support, integration support (388)

TSA1.5.4: User training in region (132)

TSA1.5.5: Site admin training in region (127)

TSA1.5.6: Regional contribution to GGUS - support teams (103)

TSA1.6: Grid Management

TSA1.6.1: OCC – CERN (60)

TSA1.6.2: Accounting coordination in region (118)

TSA1.7: Interoperation

TSA1.7.1: National and regional grid project coordination (94)

TSA1.7.2: International grid projects (39)

TSA1.8: Application<->resource provider coordination

TSA1.8.1: ROC management of resources/SLAs (119)

TSA1.8.2: OAG management (12)

TSA1.9: Application/resource provider/mw provider coord

TSA1.9.1: ROC representation in coordination (97)

TSA1.10: Network Monitoring

TSA1.10.1: Deploy network monitoring tools (84)

General Tasks (18)

Total: 5398 PM (225 FTE)

Region	Target PM (DoW)	Actual PM (expected)	Unnamed effort - PM
CERN	480	480	0
UK/I	498	498	0
France	595	623	71 [47 CSSI; 24 CNRS]
Italy	789 → 717	717	18 [INFN]
De/CH	360	360	12 [CSCS]
NE	372	372	48.5 [21.5 FOM; 13 SARA; 8 UKBH; 6 Ru-RUG]
SWE	480	480	78 [12 PIC; 48 CESGA; 6 LIP; 12 RED.Es]
SEE	744	744	221.4 [21.6 GRNET; 96 IPP, 96 UCY; 7.8 ULAKBIM]
CE	489	388	101 (tasks in WBS do not add up to target)
Ru	528	528	0
AP	135	135	20 [ASGC]
Total	5398	5325	720.3 (13.5%)

- **Highest priority:**
 - Site stability – not good!
 - Quality of new software and the effort it takes to get into production
 - Missing management functionality of middleware

- **Other issues:**
 - Some sites funded for PPS could not find effort when needed
 - This is now much better for the glite-3.0.x upgrades
 - Ensuring good communication through the ROCs to the sites
 - Often sites unaware of the schedules
 - GGUS :
 - Getting sufficient effort involved in actually doing support
 - Responsiveness and appropriateness of response

SA1 goals for EGEE-II

Key goal:

- We have a large running production infrastructure; But EGEE-II MUST take what we have now and make it:

- **Reliable**
 - Middleware components fail, error reporting is missing, ...
 - *There is an application responsibility here too - needs effort*
 - *... but ! The service has been running non-stop for > 2 years*
- **Robust**
 - Must continue to address service aspects - move away from prototypes
- **Usable**
 - It is still hard to use for many users; still too slow to introduce new VOs
- **Acceptable**
 - It must be easy to deploy in a wide variety of environments and coexist with other grid infrastructures
- **Sustainable**
 - The infrastructure must become sustainable for the long term

Reminder of what our overall goals are

SA1 Outlook

LHC VOs must achieve reliable production and analysis in 2006

- Will be making significant use of resources
- Applications must bring resources → show commitment

Consolidate and improve existing services: Focus on

- Reliability, robustness, manageability, performance, scalability, etc.
- Evolution or replacement of services driven by needs of application (or operations/security/manageability)
- ☞ TCG has key role here

- **Expand grid operations**

- Spread expertise to ROCs
- Collaboration with OSG, A-P, etc. and related projects
- Start to negotiate SLAs
- Sustainability: processes evolving, spread of expertise and tasks
- Resource sharing and negotiation – must become streamlined
 - Will need a mechanism for cost/credit for use of resources

- **System understanding (includes improving reliability)**
 - Job monitoring
 - Job and Site metrics
 - System overview
 - Etc
 - → Need a workshop/working group
- **Deployment of MPI in a more usable way**
- **Resources for new VOs**
 - Proposal under discussion – see OAG
- **Start internal review process**
 - This is in place of partner metrics which are not realistic for SA1

“Plan the migration procedure of service support for gLite in full production service more clearly with precise dates and mandates for each site, and advertise to the users well in advance.”

& comment:

“Pre-production service must not take on a life of its own...”

- **Done.**
 - *The schedule was described in the review.*
 - *Release process revised with SA3 and JRA1.*
 - *PPS now used in the way it was intended.*

“Help to establish exemplary procedures for interoperations of more divergent infrastructures and take the lead in such activities.”

- **Interoperation with OSG is a fact**
 - In use by CMS for many months
 - Operations workshop addresses several operational processes – still more work to do
 - Workflows for operational oversight and user support
 - Good ongoing relationship with OSG and many levels
- **Other efforts in progress:**
 - ARC – agreed process (+SA3 task)
 - NAREGI – workshop + GIN
 - GIN efforts – particularly around information system
 - DEISA – SA3 task for interoperability
- **LHC gives us use cases with OSG, ARC, NAREGI**
 - Not clear what the application drivers are for DEISA

- **18: “Move away from present primary dependence on particular flavours of both processors and Linux and provide support for more heterogeneous resources, including supercomputers, to allow increased collaborative adoption at major computing centres.”**
- **➔ becomes the responsibility of SA3 – see Markus’ talk**
 - *Need ETICS*
 - *Many SA3 partners involved*
 - *We need urgently to support 64-bit architectures and SLC4*

“Pursue the implementation of recommendation 18 of the second project review. Furthermore, as new infrastructures appear or the existing ones evolve, ensure interoperability according to recommendation 17 of the second review.”

- OK – see previous points.

“Improve system understanding. More specifically, investigate possible low CPU utilization, long queue lengths, and attempt to maximise system efficiency.”

- **Several ongoing efforts**

- ARDA team investigations of logs – start from experiment dashboards and follow back through LB system
 - Need to correlate with system and local logs

- **SAM, Gridview, RTM**

- Need to get the SAM tests completed – effort evaporates
 - (this is testing – and we know what that means)

- **Propose a workshop to bring together**

- SAM, WMS and LB teams, SA1, SA3, JRA2, IC (RTM) team to address this globally

Enabling Grids for E-science

“Monitor and attempt to minimize job failures.”

- **See efforts from previous point**
- **Work on gLite RB with CMS and ATLAS**
 - Need to do the same with the CE
- **Site reliability – requires better monitoring and reaction to metrics**
 - See SAM data
- **Need to train sites**
 - LCG technical day in September for LCG Tier 1s
 - Propose with HEPiX group to address Tier 2s
 - EGEE training needs to help
- **The largest source of job failures is site problems**

“Investigate how the VO resource access policies relate to the above issues and objectives.”

- Assume this means how a site allows VOs to access their resources & refers to rec. 15
- Not clear exactly what this means

“Data privacy issues may arise if the accounting data are not managed carefully. Appropriate access control policies should be explored to achieve that objective.”

- We are well aware of this issue – discussed in many fora. Problem is not at the level of the VOs – that information should be public. The problem is at the level of users.
- Problem discussed in JSPG, proposal on policy made.
 - Accounting policy, together with AUP where user agrees to certain uses of the data.
- User-level accounting will be provided by DGAS – encrypting the DN in transit, anonymise DN in the database, restrict access to certain VO members

No.	Title	Due
DSA1.1	Global Grid User Support (GGUS) Implementation Plan	1
DSA1.2	Operations Advisory Group (OAG) procedures and policy report	1
DSA1.3	Grid Services security, vulnerability and risk analysis	10
DSA1.4	Assessment of production service status	11
DSA1.5	Grid Operations cookbook	16
DSA1.6	Assessment of production grid infrastructure service status	22

No.	Title	Due
MSA1.1	Operations metrics defined	1
MSA1.2	Inventory of operations tools, procedures, & gap analysis	2
MSA1.3	Site operations policy agreement in place	5
MSA1.4	CERT teams in place – all ROCs, roles and procedures	6
MSA1.5	GGUS operational	6
MSA1.6	Security and availability policy	8
MSA1.7	Assessment of GGUS support	11
MSA1.8	Operational accounting portal	15

- **70 page document containing:**
 - Description of GGUS, and its partners (FZK, AGSC, OCC and ESC)
 - Description of the operation of GGUS
 - A work breakdown of the work into a list of 27 work items
 - Assignment of the these items to the partners
 - A list of key performance indicators for GGUS
 - Partner statements contributed by:
 - FZK, AGSC, OCC and ESC
 - Each of the 11 ROCs
 - List of exclusions from the program of work
 - Risk analysis

What is the overall plan for GGUS?

- **The EGEE-II proposal provides for the following GGUS deliverables and milestones:**
 - PM01 DSA1.1 The plan
 - PM06 MSA1.5 Milestone which shows that GGUS is operational
 - PM11 MSA1.8 Assessment of the operational state of GGUS
- **Following the assessment, it may be necessary to change some aspects of GGUS.**
- **What is not in the plan:**
 - Things which do not belong to EGEE-II/SA1 eg
 - - Details of the support units
 - - Details of the other Grid Ticketing systems
 - - Details of the Virtual Organisations
 - - Details of the longer term plan beyond 2008

- **Among the visible things which GGUS does are:**
 - Provides an infrastructure which routes requests for support from users to supporters and back
 - Creates tickets in a ticketing system
 - Provides a web site which contains interfaces to the ticketing system
 - Provides interfaces to the regional ticketing systems
 - Exchanges tickets with other Grid Ticketing systems
 - Collects statistics on tickets
 - Provides links to end user documentation
 - Provides a search tool to search the tickets
 - Provides the GOC wiki
- **GGUS has to provide a number of other services to ensure that this works:**
 - Provides a triage service to ensure accurate ticket routing
 - Trains people to work on the system
 - Organises the teams providing the triage
 - Documents and monitors the agreements on which the system is based
 - Ensures that the system is reliable and available

- **The following is the list of the more important KPIs:**
 - Number of tickets per week
 - Number of ROCs connected
 - Number of Support Units providing support
 - Number of Virtual Organisations supported
 - Number of Other Grid Systems connected
 - Mean time to solve a ticket
 - **Time to respond to tickets**
 - The maximum time to solve a ticket

- **OAG in brief**
- **Contents of deliverable**
- **Status**
 - Done / Under work / Planned
- **Issues**
- **Links**

- **High-level interface between Applications and Operations**
- **Co-chaired by one person of NA4 and one of SA1**
- **Manages procedures:**
 - To recognize new VOs
 - To define MoUs
- **Acts as a broker for resource requests**
- **Helps in identifying and handling major procedural problems between VOs and Operations**

- **History from start in EGEE-I**
- **More detailed mandate**
- **Composition**
 - Chairpersons, various representatives from NA4 and SA1
- **VO integration procedure, tools**
 - State and plans, workflow diagrams
- **Concerns**

- **New VO registration procedure in place**
 - Web form based
 - Since start in July, about 30 registration requests received, 6 not yet completed (as of 11th August 2006).
- **MoU with DILIGENT under work**
- **Still some minor streamlining needed**

- **Tools to show high level resource allocation by region and VO are planned**

- **For SA1, resource negotiation procedures have to be developed.**
 - This has to be done by region
 - Resource allocation summary tools are a pre-requisite
 - Escalation procedures in case of unsatisfied requests have to be found
- **The operation of the OAG itself has to be changed**
 - No EGAAP any longer
 - User Forum and EGEE Conference now more important for face-to-face meetings

- OAG contact: oag-chair@in2p3.fr
- OAG deliverable: <https://edms.cern.ch/document/724636>
- VO registration procedure: <https://edms.cern.ch/document/503245>
- VO registration web form:
https://cic.in2p3.fr/index.php?id=vo&subid=vo_registration

- **ROC managers' meetings**
 - Bi-weekly phone conference, + regular face-face meetings
 - Expectation that ROC managers disseminate information within the regions
- **COD (CIC on Duty) meetings**
 - Focused on operations coordination processes
- **Weekly operations meeting**
 - Address all operations issues
- **Operations workshops**
 - Twice a year → ~ all hands meetings for SA1
- **Mailing lists (many ...)**

- **Production infrastructure is in continual, significant, and increasing usage**
- **We still have problems with site reliability**
- **gLite-3.0 is now in full production**
 - But have not yet reached the point where we can retire the “old” services – this will become a support issue
- **A lot of work now needs to be focused on better understanding and tuning the system**
- **GGUS has made progress**
 - Some encouraging reports in the ops workshop
 - But still a long way to go
- **Policy issues**
 - New VOs and resource access is still problematic