



Enabling Grids for E-science

Status of EGEE Production Service

Ian Bird, CERN

SA1 Activity Leader

EGEE 1st EU Review

9-11/02/2005

www.eu-egee.org



- **Overview of the Grid Operations Service activities (SA1, SA2) – structure, successes, issues, and plans**
- **Strategy has been to**
 - have a robust certification and testing activity,
 - simplify as far as possible what is deployed, and to make that robust and useable.
 - In parallel construct the essential infrastructure needed to operate and maintain a grid infrastructure in a sustainable way.
- **Current service based on work done in LCG – culminating in the current service (“LCG-2”)**
 - Now at the point where in parallel we need to deploy and understand gLite – whilst maintaining a reliable production service.

- **Successes:**

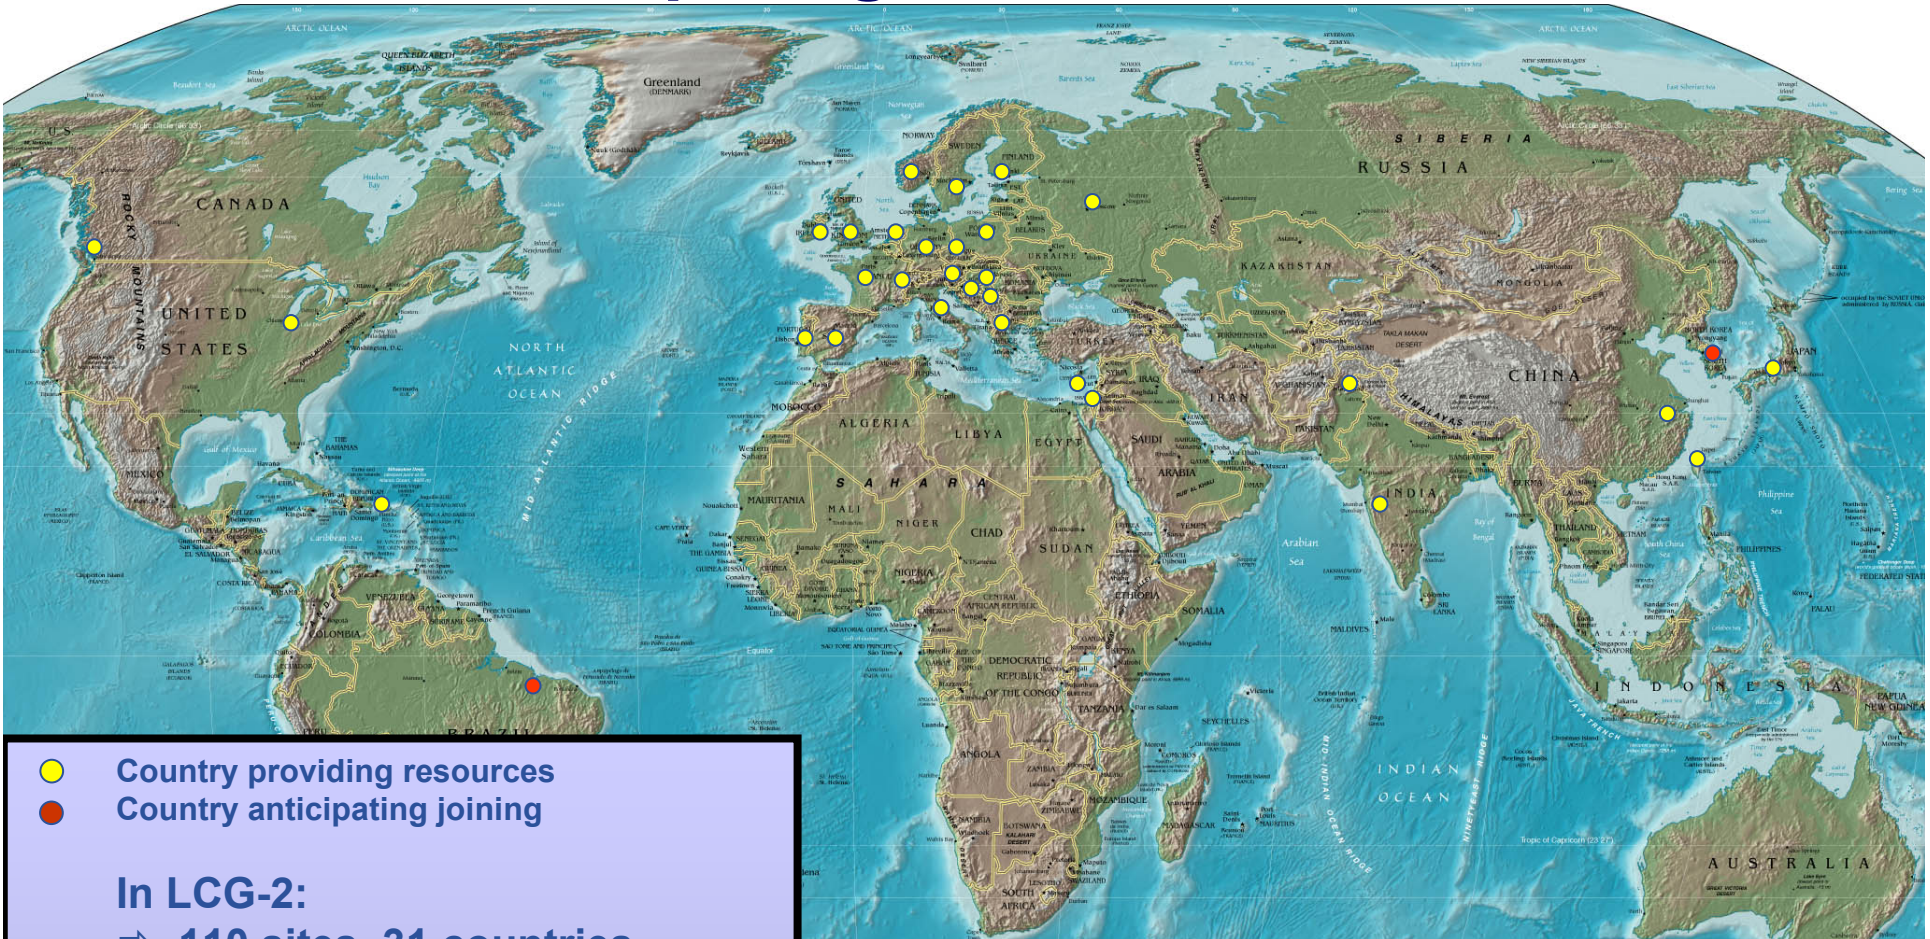
- A large operational production grid infrastructure in place and in use
 - Managed certification and deployment process in place
 - ☞ *Markus Schulz - talk*
 - Managed grid operations process in place
 - ☞ *Hélène Cordier - demo*
- Have supported extensive and intensive use by the LHC experiments during 2004 data challenges (10 months)
 - ☞ *NA4 talk*
- Now has Bio-medical community using the infrastructure, and others close

- **Issues:**

- ❖ Continue to improve the quality, reliability and efficiency of the operations
 - How to approach "24x7" global operations.
- ❖ Develop user support in order to build a trusted, reliable and usable user support infrastructure
- ❖ Introducing and deploying new VOs is too heavy weight

- **Core Infrastructure services:**
 - Operate essential grid services
- **Grid monitoring and control:**
 - Proactively monitor the operational state and performance,
 - Initiate corrective action
- **Middleware deployment and resource induction:**
 - Validate and deploy middleware releases
 - Set up operational procedures for new resources
- **Resource provider and user support:**
 - Coordinate the resolution of problems from both Resource Centres and users
 - Filter and aggregate problems, providing or obtaining solutions
- **Grid management:**
 - Coordinate Regional Operations Centres (ROC) and Core Infrastructure Centres (CIC)
 - Manage the relationships with resource providers via service-level agreements.
- **International collaboration:**
 - Drive collaboration with peer organisations in the U.S. and in Asia-Pacific
 - Ensure interoperability of grid infrastructures and services for cross-domain VO's
 - Participate in liaison and standards bodies in wider grid community

Month	Deliverable / Milestone	Item	Lead
M03	DSA1.1	Detailed execution plan for first 15 months of infrastructure operation	CERN
M06	MSA1.1	Initial pilot production grid operational 10 sites	
M06	DSA1.2	Release notes corresponding to the initial pilot Grid infrastructure operational	INFN
M09	DSA1.3	Accounting and reporting web site publicly available	CCLRC
M09	MSA1.2	First review	
M12	DSA1.4	Assessment of initial infrastructure operation and plan for next 12 months	IN2P3
M14	DSA1.5	First release of EGEE Infrastructure Planning Guide ("cook-book"),	CERN
M14	MSA1.3	Full production grid infrastructure operational 20 sites	
M14	DSA1.6	Release notes corresponding to the full production Grid infrastructure operational	CCLRC
M18	MSA1.4	Second review	
M22	DSA1.7	Updated EGEE Infrastructure Planning Guide	CERN
M24	DSA1.8	Assessment of production infrastructure operation and outline of how sustained operation of EGEE might be addressed.	IN2P3
M24	MSA1.5	Third review and expanded production grid operational 50 sites	
M24	DSA1.9	Release notes corresponding to expanded production Grid infrastructure operational	INFN




- Country providing resources
- Country anticipating joining

In LCG-2:

- ⇒ 110 sites, 31 countries
- ⇒ >10,000 cpu
- ⇒ ~5 PB storage

Includes non-EGEE sites:

- 10 countries
- 18 sites



- VOs and users

Metrics	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Details
Number of supported VOs	8	9	10	See details
Number of associated VO	.	40	44	See details
Supported VOs not primarily from physics	.	5	5	Biomed, ESR (Earth Sciences), Compchem (Chemistry), Magic (Astronomy), Egeod (Geo-Physics)
Number of users in supported VOs	.	.	497(*)	See details .
Number of users in associated VOs	See details . Accurate numbers will be provided in the next QR.
Number of disciplines	.	5	6	Chemistry, Astronomy, Physics, Earth Sciences, BioMed, Geo-Physics See disciplines for supported and associated VOs
Number of experiments from physics	.	7	7	LHC: ALICE , ATLAS , CMS , LHCb , More details Non-LHC: DO , Barbar , CDF , more details
Number of deployed applications not primarily from physics approved by EGAAP	<p style="text-align: center; font-size: 2em; color: blue;">● To be updated</p>								
Applications deployed for testing on GILDA									
Number of applications submitted to EGAAP	.	4	4+6	See details
Number of countries	.	26	27	See details

Fed.	Annex 1 Expectation		Status at	Status at
	PM1	PM15	PM6	PM9
CERN	900	1800	956	940
UK	100	2200	2132	2415
FR	400	895	160	244
IT	553	679	1836	1337
SE	146	322	108	130
SW	250	250	408	390
CE	385	730	356	327
NE	200	2000	348	364
DE-CH	100	400	910	1161
RU	50	152	169	156
Totals	3084	9428	7383	7464

• To be updated & include actual situation,
 • No. not in Europe

Metrics	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Details
OMC		1	1						See details
CIC		5	5						See details
ROC		11	11						See details
RCs		<u>67</u>	<u>75</u>						

- **Mechanics:**

- The recipe is straightforward and clear
- But, this is a heavy weight process and must be improved
- Requires a lot of configuration changes by a site
 - Often leads to problems

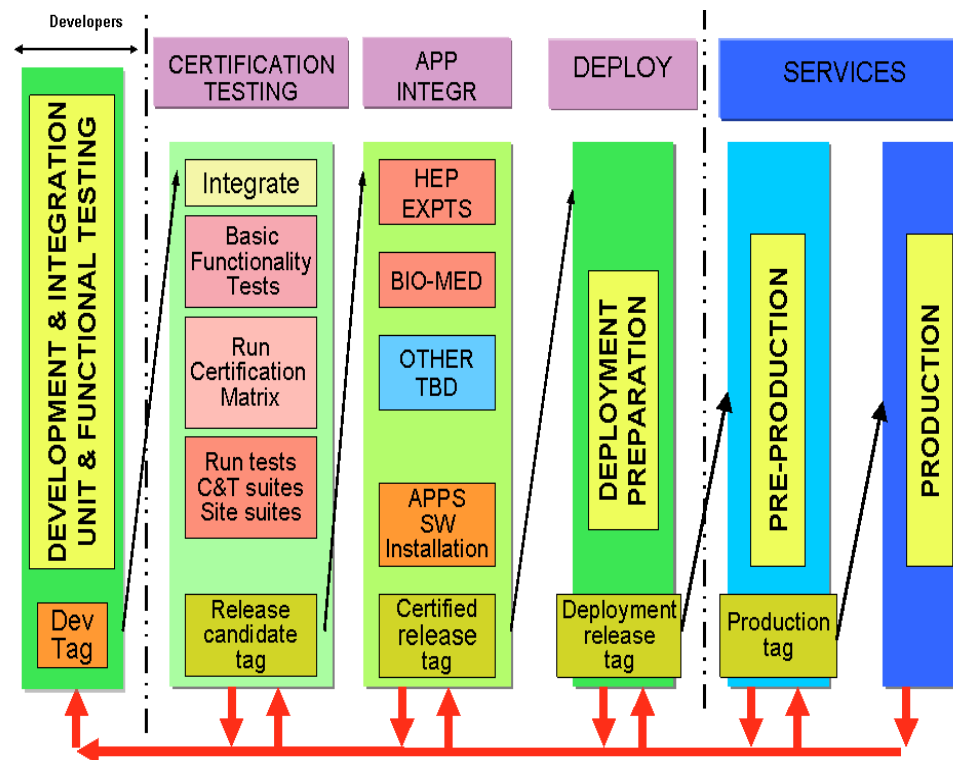
- **Policy:**

- Joint group of SA1/NA4 (called OAG in the TA)
- Members are the application representatives and the ROC managers; chaired by NA4
 - Mandate
 - *Understand application resource requirements*
 - *Negotiate those resources within the federations – the ROC manager is responsible to make the negotiation*
 - NB. A site is often funded for specific applications – it is by and large NOT the case that any application is entitled to run anywhere
 - *But let's demonstrate the value of being able to do that ...*

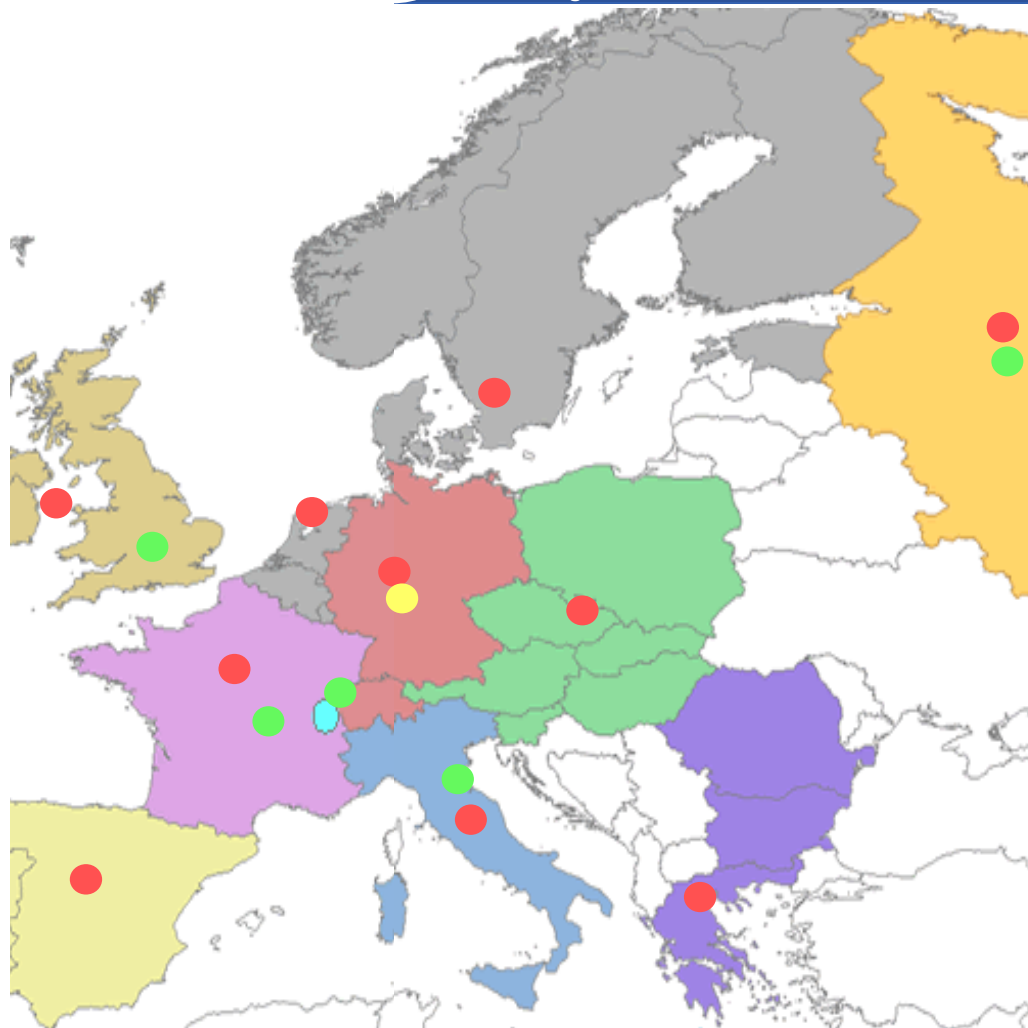
- **Evolution through 2003/2004**
 - Focus has been on making these reliable and robust
 - Basic functionality and reliability rather than additional functionality
 - Respond to needs of users, admins, operators
- **The software stack is the following:**
 - Virtual Data Toolkit
 - Globus (2.4.x), Condor, etc
 - EDG developed higher-level components
 - Workload management (RB, L&B, etc)
 - Replica Location Service (single central catalog), replica management tools
 - R-GMA as accounting and monitoring framework
 - VOMS being deployed now
 - Operations team re-worked components:
 - Information system: MDS GRIS/GIIS → BDII
 - edg-rm tools replaced and augmented as lcg-utils
 - Developments on:
 - *Disk pool managers (dCache, DPM)*
 - *Catalogue*
 - Other tools as required:
 - e.g. GridIce - DataTag

- Maintenance agreements with:
 - VDT team (inc Globus support)
 - WLM, VOMS – Italy
 - DM – CERN

- **Key point – a certification process is essential**
 - However, it is expensive (people, resources, time)
 - But, this is the only way to deliver production quality services
 - LCG-2 was built from a wide variety of “research” quality code
 - Lots of good ideas, but little attention to the boring stuff
 - Building a reliable distributed system is hard –
 - Must plan for failure, must provide fail-over of services, etc
 - Integrating components from different projects is also difficult
 - Lack of common standards for logging, error recovery, etc
- → **Markus Schulz talk**



- **The EGEE production grid service is quite stable**
 - The services are quite reliable
 - Remaining instabilities in the IS are being addressed
 - Sensitivity to site management
 - Underlying problems in (for example gridftp) must be addressed (reliable file transfer service)
- **The biggest problem is stability of sites**
 - Configuration problems due to complexity of the middleware
 - Fabric management at less experienced sites
- **Job efficiency is not high, unless**
 - Operations/Applications select stable sites (BDII allows a application-specific view)
- **Operations workshop last November to address this**
 - Fabric management working group – write fabric management cookbook
 - Tighten operations control of the grid – escalate and remove bad sites



Operations Management Centre (OMC):

- At CERN – coordination etc

Core Infrastructure Centres (CIC)


- Manage daily grid operations – oversight, troubleshooting
- Run essential infrastructure services
- Provide 2nd level support to ROCs
- UK/I, Fr, It, CERN, + Russia (M12)
- Taipei also run a CIC

Regional Operations Centres (ROC)

- Act as front-line support for user and operations issues
- Provide local knowledge and adaptations
- One in each region – many distributed

User Support Centre (GGUS)

- In FZK – manage PTS – provide single point of contact (service desk)
- Not foreseen as such in TA, but need is clear



• RC = Resource Centre

- The *grid* is flat, but
- **Hierarchy of responsibility**
 - Essential to scale the operation
- **CICs act as a single Operations Centre**
 - Operational oversight (*grid operator*) responsibility
 - rotates weekly between CICs
 - Report problems to ROC/RC
 - ROC is *responsible* for ensuring problem is resolved
 - ROC oversees regional RCs
- **ROCs responsible for organising the operations in a region**
 - Coordinate deployment of middleware, etc
- **CERN coordinates sites not associated with a ROC**

- **CIC-on-duty**

- Responsibility rotates through CIC's – one week at a time
- Manage daily operations – oversee and ensure
 - Problems from all sources are tracked (entered into PTS)
 - Problems are followed up
 - CIC-on-duty hands over responsibility for problems
- Hand-over in weekly operations meeting

- **Daily operations:**

- Checklist
- Various problem sources: monitors, maps, direct problem reports

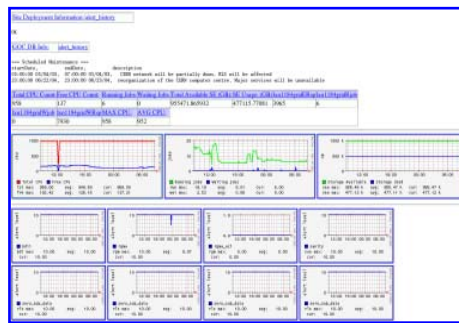
- **Next step:**

- Continue to develop tools to generate automated alarms and actions


Enabling Grids for E-science



GIIS Monitor



GIIS Monitor graphs




Site Functional Tests

Variety of monitoring tools are in daily use.
 Some of these you see on the live display monitors.
 More details in SA1 demo.



GOC Data



Job Monitor



Gridce - VO view



Gridce - fabric view



Certificate Lifetime Monitor

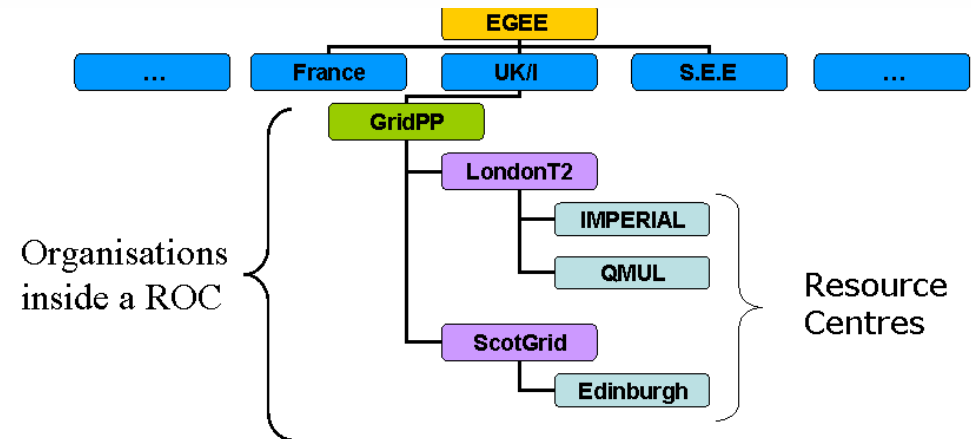

Note: Those thumbnails are links and are clickable.

- **Need service level definitions**
 - What a site supports (apps, software, MPI, compilers, etc)
 - Levels of support (# admins, hrs/day, on-call, operators...)
 - Response time to problems
 - Agreement (or not) that remote control is possible (conditions)
- **Sites sign-off on responsibilities/charter/SLD**
- **Publish sites as bad in info system**
 - Based on unbiased checklist (written by CICs)
 - Consistently bad sites → escalate to political level GDB/PMB
- **Small/bad sites**
 - “Force” sites to follow upgrades
 - Remote management of services
 - Remote fabric monitoring (GridICE etc)

- **How to move towards a 24x7-like global support:**
 - Separate security (urgent issues) from general support
 - Distributed CIC provides “24x7” by using EGEE, Taipei, (America/Canada?)
 - Real 24x7 coverage only at CERN and large centres (CIC-centres)
 - Or other specific crucial services that justify cost
 - Loss of capacity – vs damage
 - Classify what are 24x7 problems
 - Direct user support not needed for 24x7
 - Massive failures should be picked by operations tools
- **Having an operating production infrastructure should not mean having staff on shift everywhere**
 - “best-effort” support
 - The infrastructure (and applications) must adapt to failures

- Accounting at the moment is “after the fact”
 - The most important way to determine how many resources were consumed by each VO (and potentially each user)
 - No attempt to establish or impose quotas
 - But of course, each site can and does do so
 - Not a trivial problem – jobs should not go to a site where they have no resource, but a modern batch system cannot give a definitive reply

Accounting Flow Diagram



Accounting

- > Accounting Home
- Accounting Plots**
- > CIC View
- > ROC View
- > Country View
- > Custom Query
- > Site View
- General**
- > Contact the GOC accounting team



GOC Accounting Services

Latest News



- **Dec-08-2004:**
The accounting RPMs packaged in LCG 2_3_0 (version 3.4.37) have been recently updated! The latest release (3.4.38) is found [here](#).
- **Dec-02-2004:**
The Accounting Service is live.

Database Statistics

Total Number of Records in Database



Total number of distinct sites reporting to GOC



Accounting menu may be used to select different views of the data

- **Operational Security team in place**

- EGEE security officer, ROC security contacts
- Concentrate on 3 activities:
 - Incident response

• **Security group and work was started in LCG – was from the start a cross-grid activity.**

• **Much already in place at start of EGEE: usage policy, registration process and infrastructure, etc.**

• **We regard it as crucial that this activity remains broader than just EGEE**

handling and response to cyber security incidents on Grids

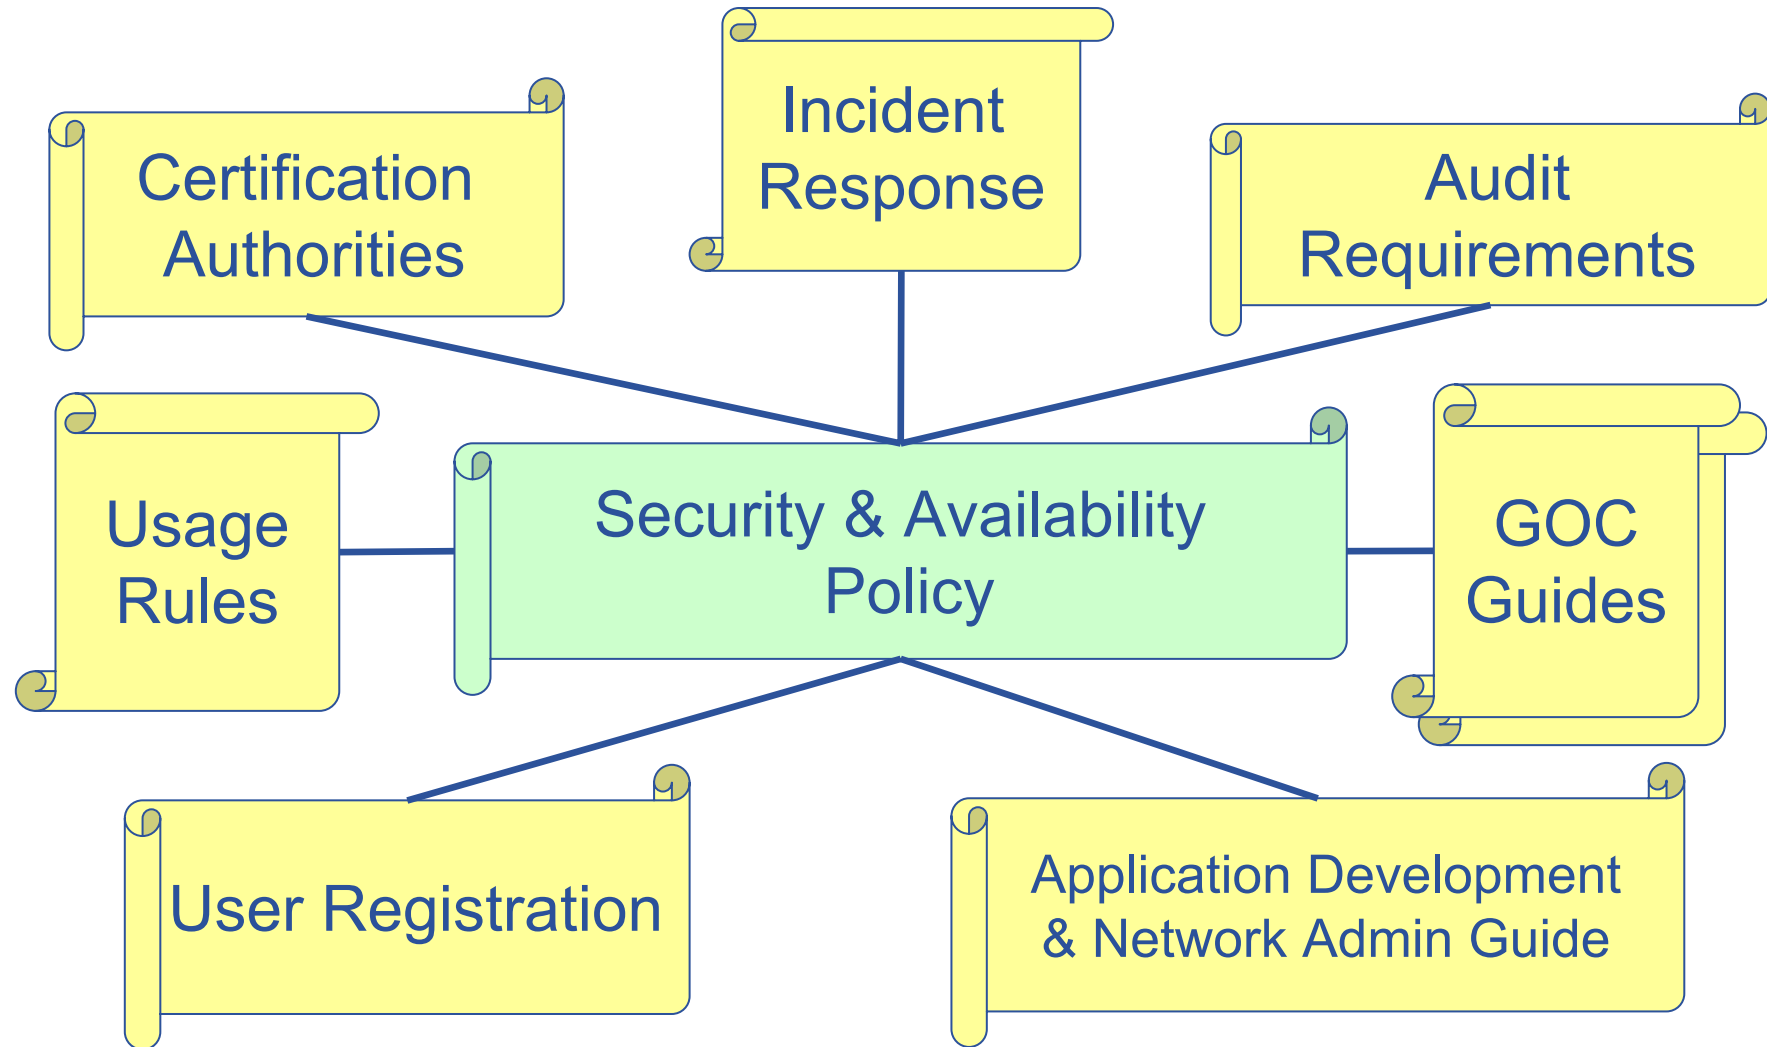
- Basic framework for incident definition and handling

- **Site registration process in draft**

- Part of basic SLA

- **CA Operations**

- EUGridPMA – best practice, minimum standards, etc.
- More and more CAs appearing



<http://cern.ch/proj-lcg-security/documents.html>

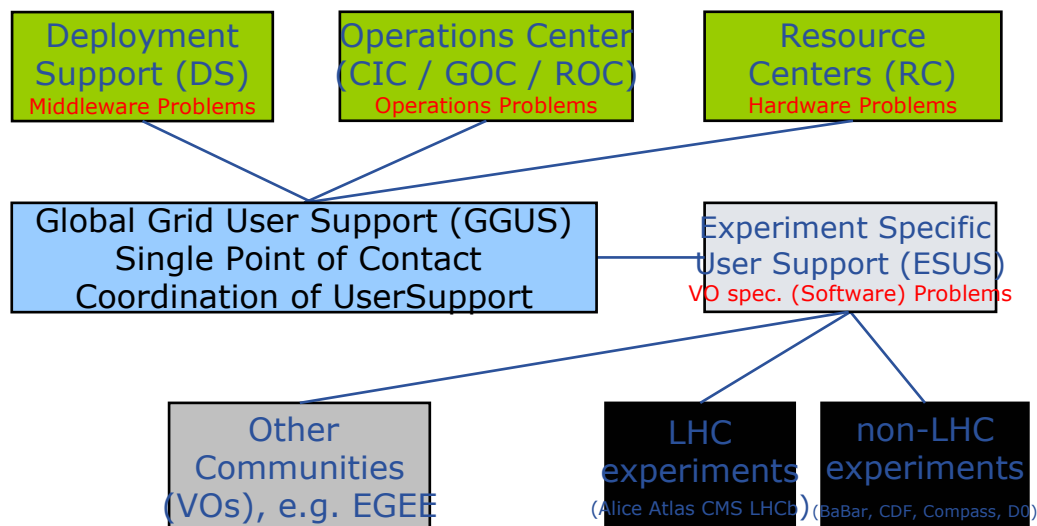
We have found that user support has 2 distinct aspects:

- User support –**

- Call centre/helpdesk
- Coordinated through GGUS
- ROCs as front-line
- Task force in place to improve the service

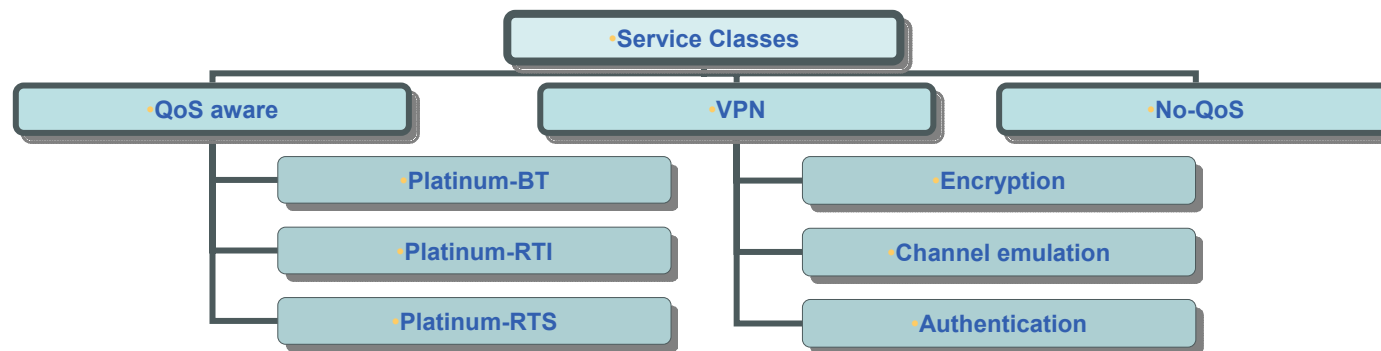
- VO Support**

- Was an oversight in the project and is not really provisioned
- In LCG we have a team (5 FTE):
 - Help apps integrate with m/w
 - Direct 1:1 support
 - Understanding of needs
 - Act as advocate for app
- This is really missing for the other apps – adaptation to the grid environment takes expertise




- **National Grids within EGEE**
 - The large national grid infrastructures in EGEE regions are becoming integrated into the overall service:
 - Italy – Grid.IT sites are part of EGEE
 - UK/I – National Grid Service sites are part of EGEE
 - Nordic countries – Some sites run EGEE in parallel with NorduGrid
 - [+ SEE-Grid + EELA]
- **Strong relationship with Asia-Pacific**
 - Taipei acts as CIC and hopefully will become a ROC
- **External Grids**
 - Most important are Grid3 (→ Open Science Grid) in USA and the Canadian Grid efforts (WestGrid and GridCanada)
 - OSG and EGEE use same base sw stack – we have demonstrated job interoperability in both directions
 - *Operations and security teams have much in common – proposing specific joint activities*
 - Canada – at Triumf a gateway from EGEE to Canadian resources has been built and used in production
- **This momentum has to maintained as we move to the next generation of middleware**

- **Technical Network Liaison Committee setup:**
 - To provide an efficient place to deal with “practical” issues of interface between NRENs and EGEE (Network SLAs, Network Services),
 - 8 members: EGEE (SA2, SA1, JRA4), GEANT/NRENs (DANTE, DFN, GARR, GRNET), CERN,
 - 2 meetings in Cork and Den Haag.
- **First survey of network requirements complete**
 - A SA2-JRA4 workgroup has gathered 36 requirements, available in the EGEE requirements database.
 - Three main requirement classes (operational, flow control, network characteristics) allow the specification of a minimum level of services (SLRs).
- **First service classes identified**
 - « User oriented » service classes, not « network classical classification».



- **European network services survey, 43 NRENs concerned**
 - Questionnaire sent to NRENs,
 - Data extracted from the TERENA compendium.



- **QoS experimentation → A real network QoS use case in EGEE**
 - Application: GATE (Geant4 Application for Tomographic Emission),
 - NRENs involved: Renater, RedIris, GEANT,
 - Aim: To have a better approach for the SLAs processing, to ask for network requirements to the middleware.

- **Initial model for network service usage (M9)**
 - A mapping of the EGEE services classes in the NRENs services classes,
 - *Platinum-RTI and Platinum RTS in Premium IP (PIP) service,*
 - *Platinum-BT in the Best Effort Service or LBE service,*
 - *No available solution for VPN Encryption and Authentication,*
 - *For channel emulation, the service is only available in some parts of the networks.*
 - A generic model for network resource management taking into account different provisioning mechanisms.
 - A SLS template which will be the technical part of the SLA.

- **For the next period (M10-M24)**
 - SLAs definition, implementation:
 - Based on the previous works and the responses from EGEE and GN2 to some open issues (procedures, demarcation point ...)
 - Definition in cooperation with GN2.
 - Operational interface between EGEE and GEANT/NRENs
 - *SLA agreements processing, SLA monitoring.*
 - *Trouble Ticket system & reporting procedures*
 - To have a theoretical schema approved by the partners (M12),
 - To implement the operational model in order to have a mature network operational interface.

- **Milestones**

- MSA1.3 (M14) Full production grid infrastructure operational
 - 20 sites, using re-engineered middleware
- MSA1.4 (M18) Second project review
- MSA1.5 (M24) Expanded production grid operational
 - 50 sites

- **Deliverables**

- DSA1.4 (M12) Assessment of operation of 1st 12 months
- DSA1.5 (M14) First release of “cook-book”
- DSA1.6 (M14) Release notes corresponding to MSA1.3
- DSA1.7 (M22) Second edition of “cook-book”
- DSA1.8 (M24) Assessment of production operation
 - Include thoughts on how to make the infrastructure sustainable
- DSA1.9 (M24) Release notes corresponding to MSA1.5

- **Changes wrt TA**

- No significant change

- TBD ... “We have done a lot – a lot more still to do”