



Enabling Grids for E-science

SA1

Ian Bird
SA1 Activity Leader
CERN IT Department

EGEE Final Review

23rd – 24th May 2006

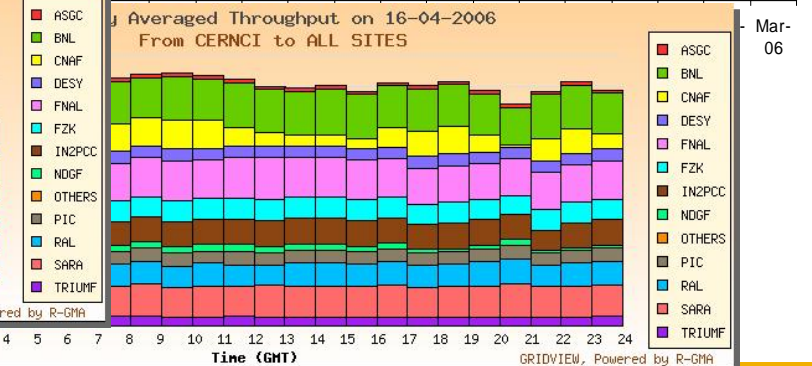


www.eu-egee.org



Information Society
and Media



- Recommendations from intermediate focused review
- Highlights of last 3 months of the project
- Summary of SA1 achievements and open issues



“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.”

- Early set up of TCG;
 - forum for agreeing schedules across the technical and application activities.
 - Schedule proposed and agreed for 2006 – see next slide

Deployment schedule for 2006

- **Deliver and deploy LCG-2.7.0 end January 2006**
 - Bug fixes, patches, etc. accumulated since last major
- **Quickfixes, security patches**
 - May be produced at any time, deployed with agreement of TCG
- **Client tools**

In general we try to move away from big-bang releases:

- Focus on service/component upgrades where possible
- Check-point releases to consolidate changes and to provide new sites a starting point
- See this more like a Linux distribution – major releases with continual component updates, security patches, etc.

Evolutionary from deployment point of view – will not be a big-bang change of production service

- Schedule driven by LCG service challenges
- To demonstrate functionality or provide new facilities
- Usually need by-hand installation
- **Foresee second major “release” on October/November timescale**
 - Added functionality – driven by apps via TCG

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

- **Number of initiatives: OSG, Nordugrid, GIN**
 - Security
 - Storage interfaces
 - Information systems
 - Job submission
 - Monitoring
 - Common operations policies
 - Operations workshops
- **Joint operations:**
 - WLCG is a strong driver – bring together EGEE and OSG grid operations
- **Relations with other projects:**
 - EU related infrastructure projects
 - OSG, NorduGrid – collaboration to provide WLCG service
 - NAREGI – exploratory workshop; support for Japanese WLCG sites
 - DEISA, ARC – specific tasks in EGEE-II to work on interoperability
 - Grid Interoperability Now (GIN) group – many projects

“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.”

- **Current porting status:**

- Several ports to other architectures: IA64, several Linux flavours. Available a few months after main release;
- Done by remote partners; outside of main build and integration system

- **Future:**

- Important to have several important ports close to or part of main integration and testing;
- Include 64-bit cleanliness as part of build test – will flag as failure
- Move to ETICS to provide distributed build system to support many platforms; helps tie porting partners into central process
 - Partner interested in a particular port can provide build and test hardware and ETICS can help integrate this into the process
- TCG should agree a reasonable/realistic set of standard primary platforms to be provided as part of base release
 - E.g. SL4 + Debian on 32 and 64 bit
 - Other ports can be asynchronous and should be certified by partners providing resources



Enabling Grids for E-scienceE

SA1 Highlights

www.eu-egee.org





INFSO-RI-508833



EGEE Grid Sites : Q1 2006

EGEE:
Steady growth over the lifetime of the project

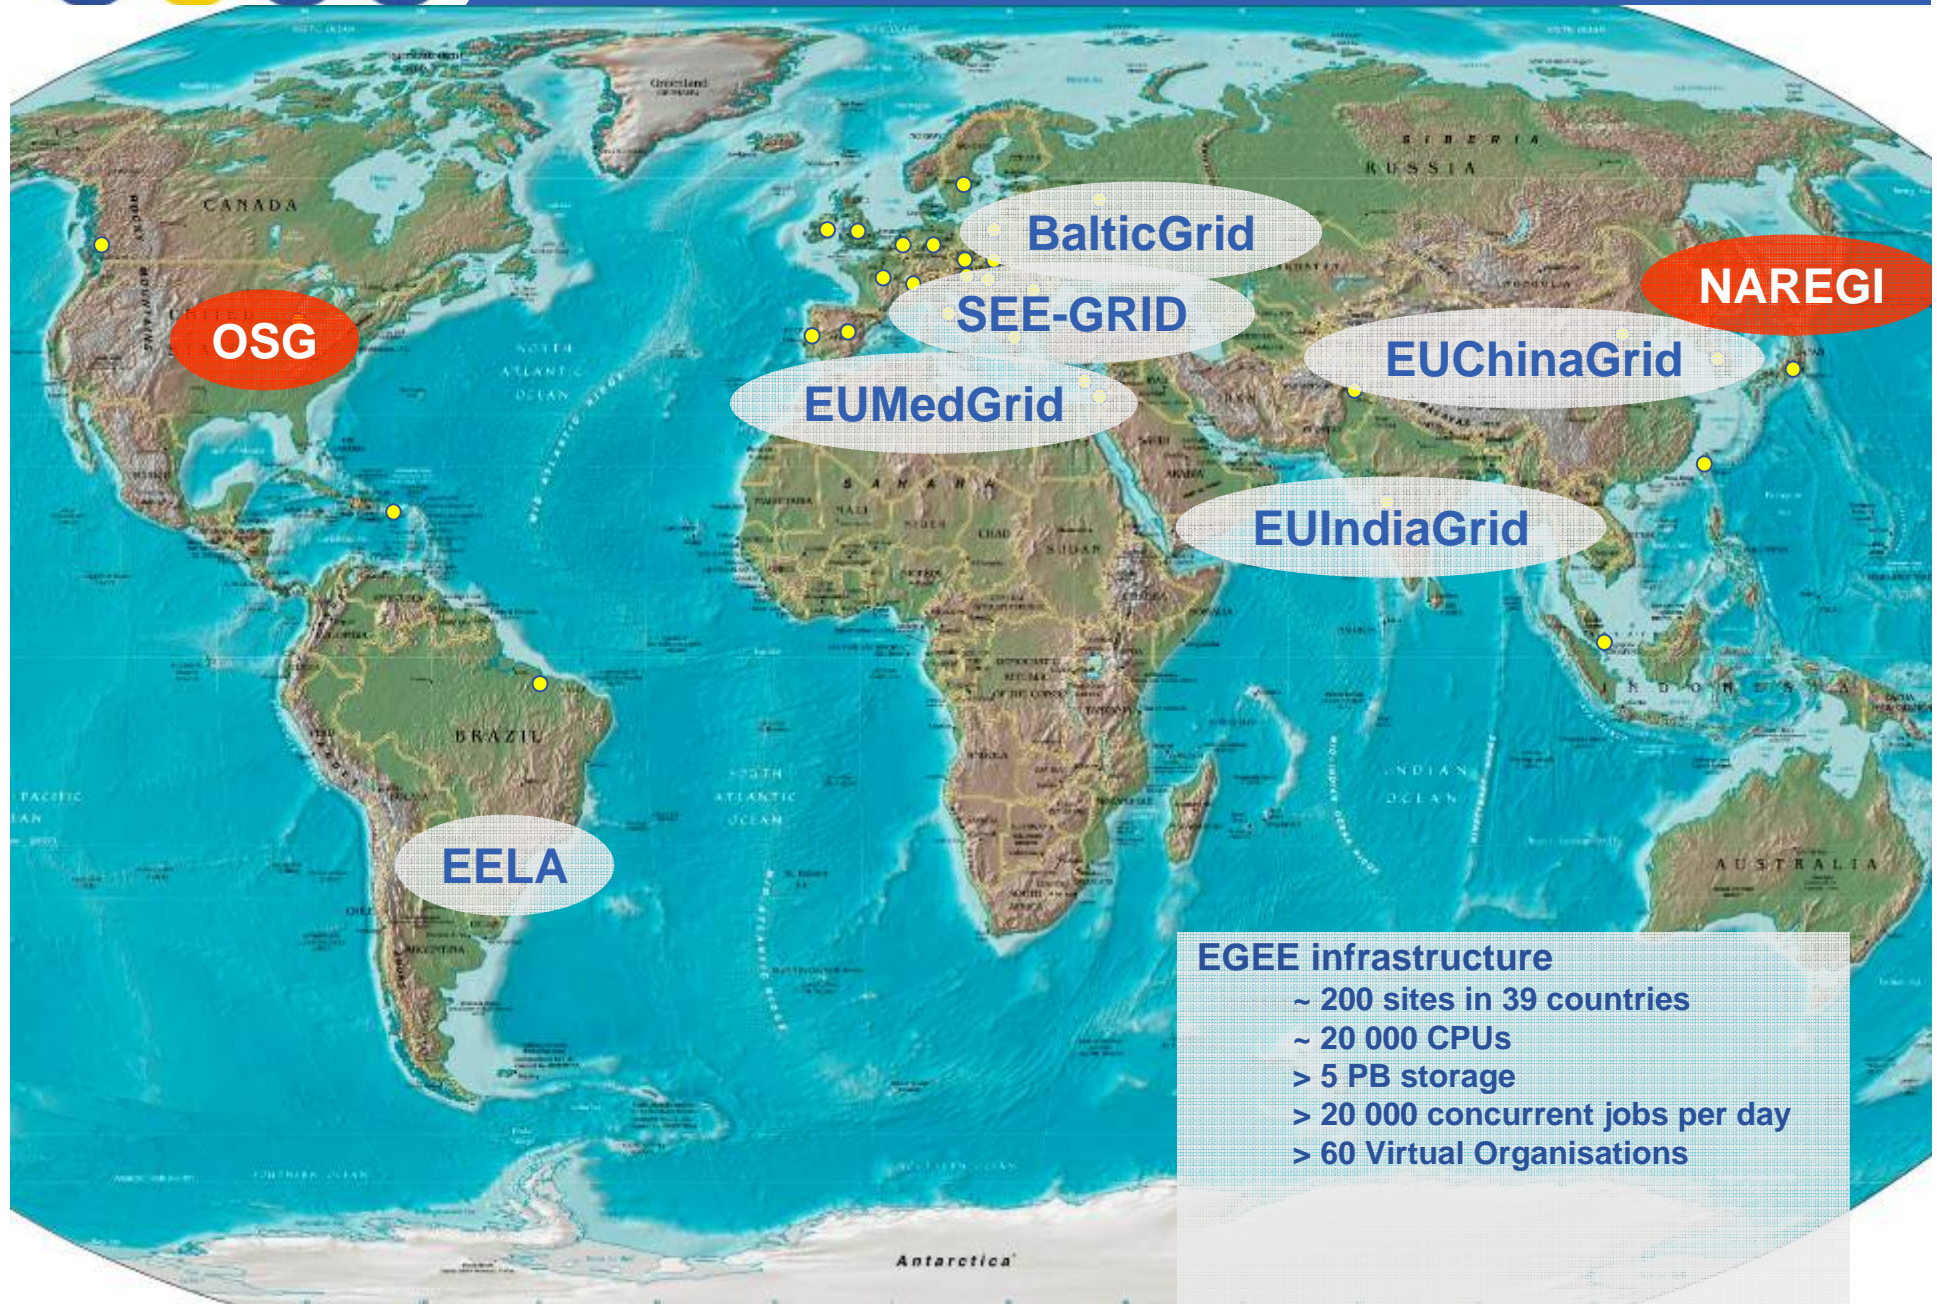


EGEE:
 > 180 sites, 40 countries
 > 24,000 processors,
 ~ 5 PB storage

country	sites	country	sites	country	sites
Austria	2	India	2	Russia	12
Belgium	3	Ireland	15	Serbia	1
Bulgaria	4	Israel	3	Singapore	1
Canada	7	Italy	25	Slovakia	4
China	3	Japan	1	Slovenia	1
Croatia	1	Korea	1	Spain	13
Cyprus	1	Netherlands	3	Sweden	4
Czech Republic	2	Macedonia	1	Switzerland	1
Denmark	1	Pakistan	2	Taipei	4
France	8	Poland	5	Turkey	1
Germany	10	Portugal	1	UK	22
Greece	6	Puerto Rico	1	USA	4
Hungary	1	Romania	1	CERN	1





A global, federated e-Infrastructure



EGEE infrastructure



- ~ 200 sites in 39 countries
- ~ 20 000 CPUs
- > 5 PB storage
- > 20 000 concurrent jobs per day
- > 60 Virtual Organisations



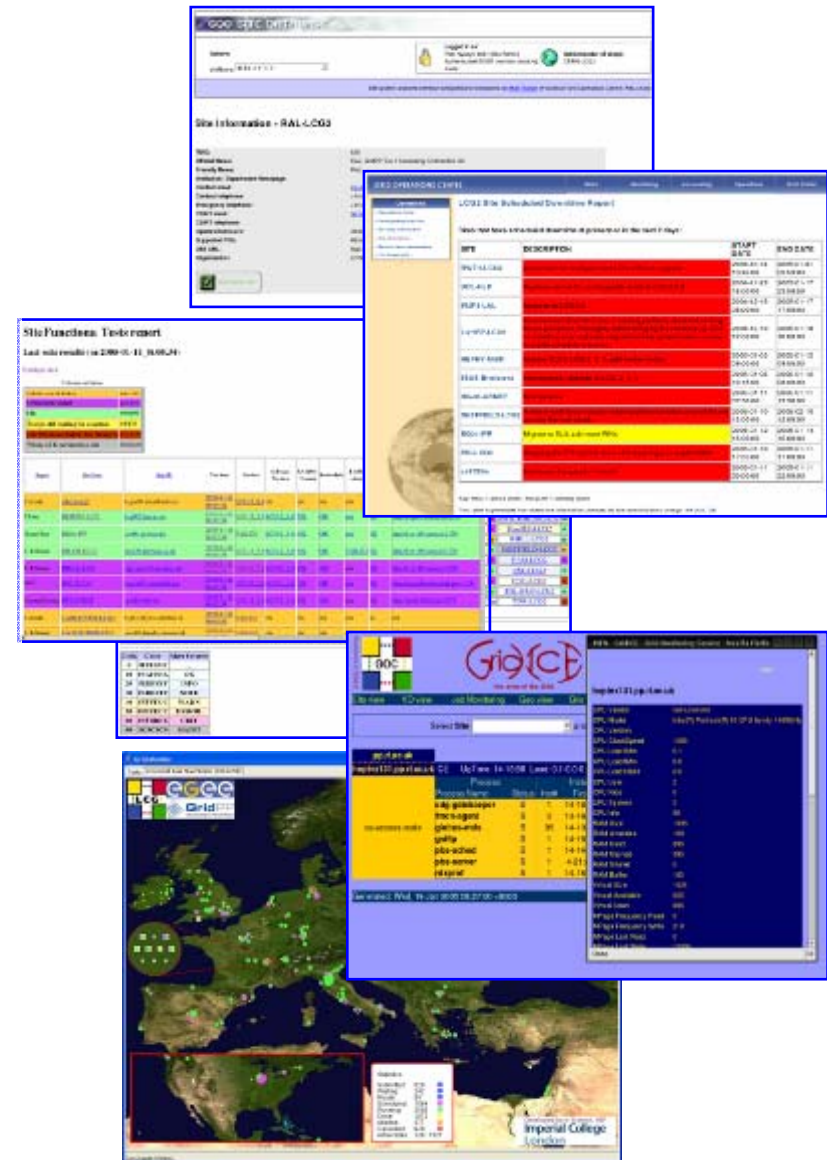
Sustained & regular workloads of 20K jobs/day

- spread across full infrastructure
- doubling in last 6 months – no effect on operations

Massive data transfers > 1.5 GB/s



- **Grid operator on duty**
 - 6 teams working in weekly rotation
 - CERN, IN2P3, INFN, UK/I, Ru, Taipei
 - Crucial in improving site stability and management
 - Expanding to all ROCs in EGEE-II
- **Operations coordination**
 - Weekly operations meetings
 - Regular ROC managers meetings
 - Series of EGEE Operations Workshops
 - Nov 04, May 05, Sep 05, June 06
- **Geographically distributed responsibility for operations:**
 - There is no “central” operation
 - **Tools are developed/hosted at different sites:**
 - GOC DB (RAL), SFT (CERN), GStat (Taipei), CIC Portal (Lyon)
- **Procedures described in Operations Manual**
 - Introducing new sites
 - Site downtime scheduling
 - Suspending a site
 - Escalation procedures
 - etc



The collage displays several key operational tools:

- GOC DB (Grid Operations Database):** A web interface showing site information for RAL-CC03, including a table with columns for SITE, DESCRIPTION, START DATE, and END DATE. The table lists various sites with their respective descriptions and dates.
- Site Functions: Tests report:** A table showing test results for various sites, with columns for Site, Status, and Test results.
- GStat (Grid Statistics):** A web interface showing a list of sites and their operational status.
- CIC Portal (CERN Information Center):** A web interface for site management and information.
- Map:** A map of Europe showing the geographical distribution of EGEE sites across different countries.

- **Site Functional Tests (SFT)**
 - Framework to test (sample) services at all sites
 - Shows results matrix
 - Detailed test log available for troubleshooting and debugging
 - History of individual tests is kept
 - Can include VO-specific tests (e.g. sw environment)
 - Normally >80% of sites pass SFTs
 - NB of 180 sites, some are not well managed

- **Very important in stabilising sites:**
 - Apps use only good sites
 - Bad sites are automatically excluded
 - Sites work hard to fix problems

Colours definition			ver	Software Version (WN)
SD	Scheduled downtime	#a3a3a3	wn	WN host name
JL	Job list match failed	#aab3ff	ca	CA certs version
JS	Job submission failed	#f4876b	crl	CRL timestamp test
CT	Critical tests failed	#f9d48e	rm	Replica Management
NT	Non-critical tests failed	#f2f98e	votag	VO Tag management
OK	OK	#b2f98e	js	Job submission
			bi	BrokerInfo

Test summary						
	SD	JL	JS	CT	OK	total
dteam	15	12	4	6	139	176
lhcb	15	81	5	35	39	175

Extending to service availability:

- measure availability by service, site, VO
- each service has associated service class defining required availability (Critical, highly available, etc.)

First approach to SLA

Use to generate alarms

- generate trouble tickets
- call out support staff

St.	Site Name	Site CE	VO dteam											VO lhcb					
			St.	js	ver	wn	ca	rgma	bi	cs	rm	votag	swdir	crl	St.	js	dirac-test		
AsiaPacific																			
CT	INDIACMS-TIFR	ce.indiacms.res.in	CT	O	2	6	0	I	O	O	O	O	X	O	O	!!!	JL	X	??
OK	TW-NCUHEP	grid01.phy.ncu.edu.tw	OK	O	2	6	0	I	O	O	O	O	O	O	O	!!!	JL	X	??
OK	TOKYO-LCG2	dgce0.icepp.jp	OK	O	2	4	0	I	O	O	O	O	O	O	O	!!!	JL	X	??
OK	Taiwan-LCG2	lcg00125.grid.sinica.edu.tw	OK	O	2	6	0	I	O	O	O	O	O	O	O	!!!	JL	X	??
OK	Taiwan-IPAS-LCG2	testbed001.phys.sinica.edu.tw	OK	O	2	6	0	I	O	O	O	O	O	O	O	!!!	JL	X	??
OK	GOG-Singapore	melon.ngpp.ngp.org.sg	OK	O	2	6	0	I	O	O	O	O	O	O	O	!!!	JL	X	??
OK	Taiwan-NCUCC-LCG2	ce.cc.ncu.edu.tw	OK	O	2	6	0	I	O	O	O	O	O	O	O	!!!	OK	O	O
OK	LCG_KNU	cluster50.knu.ac.kr	OK	O	2	5	0	I	O	O	O	O	O	O	O	!!!	CT	O	!!!
BNL																			
SD	BNL-LCG2	lcg-cc01.usatlas.bnl.gov	SD	X	??	??	??	??	??	??	??	??	??	??	??	??	SD	X	??
Canada																			
JL	TORONTO-LCG2	bigmac-lcg-ce.physics.utoronto.ca	JL	X	2	6	0	I	O	O	O	O	O	W	O	!!!	OK	O	O
11.	SD	CARLETONU-LCG2	lcg02.physics.carleton.ca	SD	X	??	??	??	??	??	??	??	??	??	??	??	SD	X	??
12.	OK	TRIUMF-LCG2	lgce01.triumf.ca	OK	O	2	6	0	I	O	O	O	O	O	O	O	OK	O	O
13.	OK	Umontreal-LCG2	lcg-ce.lps.umontreal.ca	OK	O	2	6	0	I	O	O	O	O	W	O	!!!	OK	O	O

- **Terminology:**

- EGEE deploys a middleware distribution
 - Drawn from various middleware products, stacks, etc.
 - Do not confuse the *distribution* with development projects or with software packages
 - Count on 6 months from software developer “release” to production deployment
- The EGEE distribution:
 - Current production version labelled: LCG-2.7.0
 - Next version labelled: gLite-3.0
 - Name change to hopefully *reduce* confusion



- **EGEE distribution contents:**

- ❖ **LCG-2.7.0:**

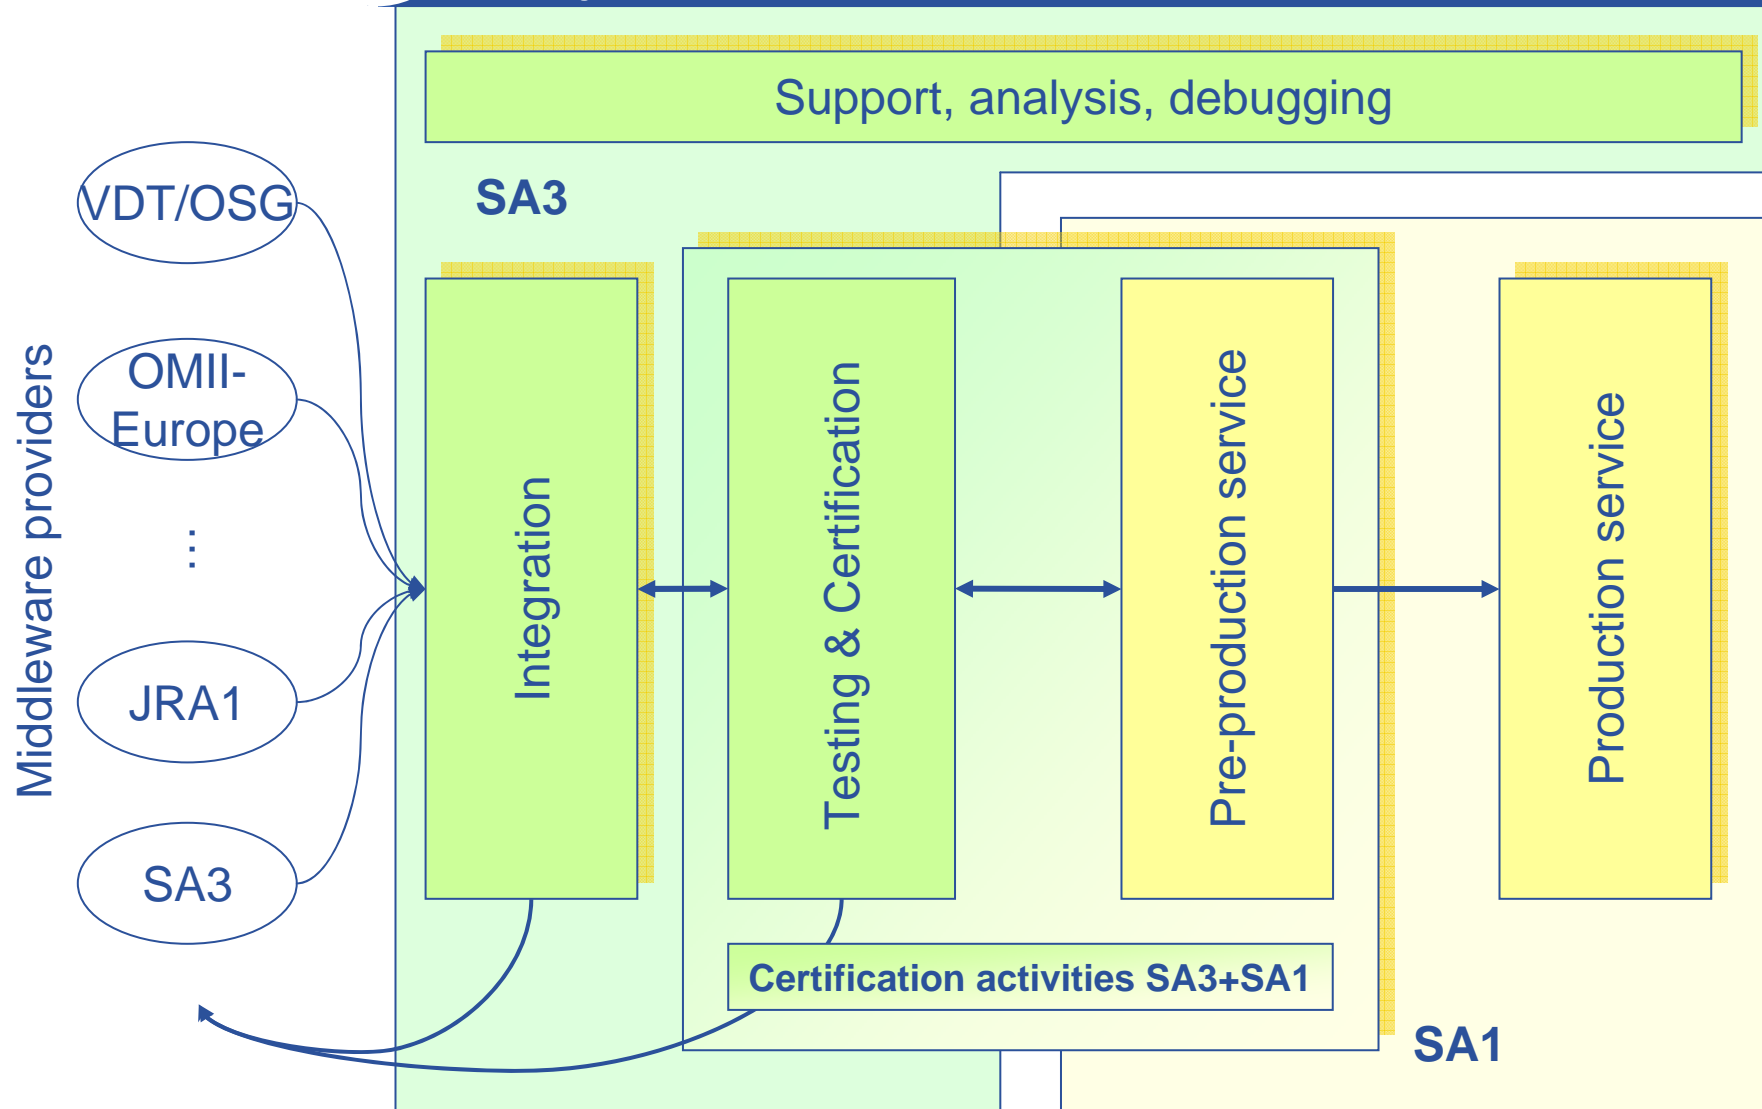
- VDT – packaging Globus 2.4, Condor, MyProxy
- EDG workload management
- LCG components:
 - BDII (info sys),
 - catalogue (LFC),
 - DPM, data management libraries and CLI tools
 - monitoring tools
- gLite: R-GMA, VOMS, FTS

evolution →

- ❖ **gLite-3.0:**

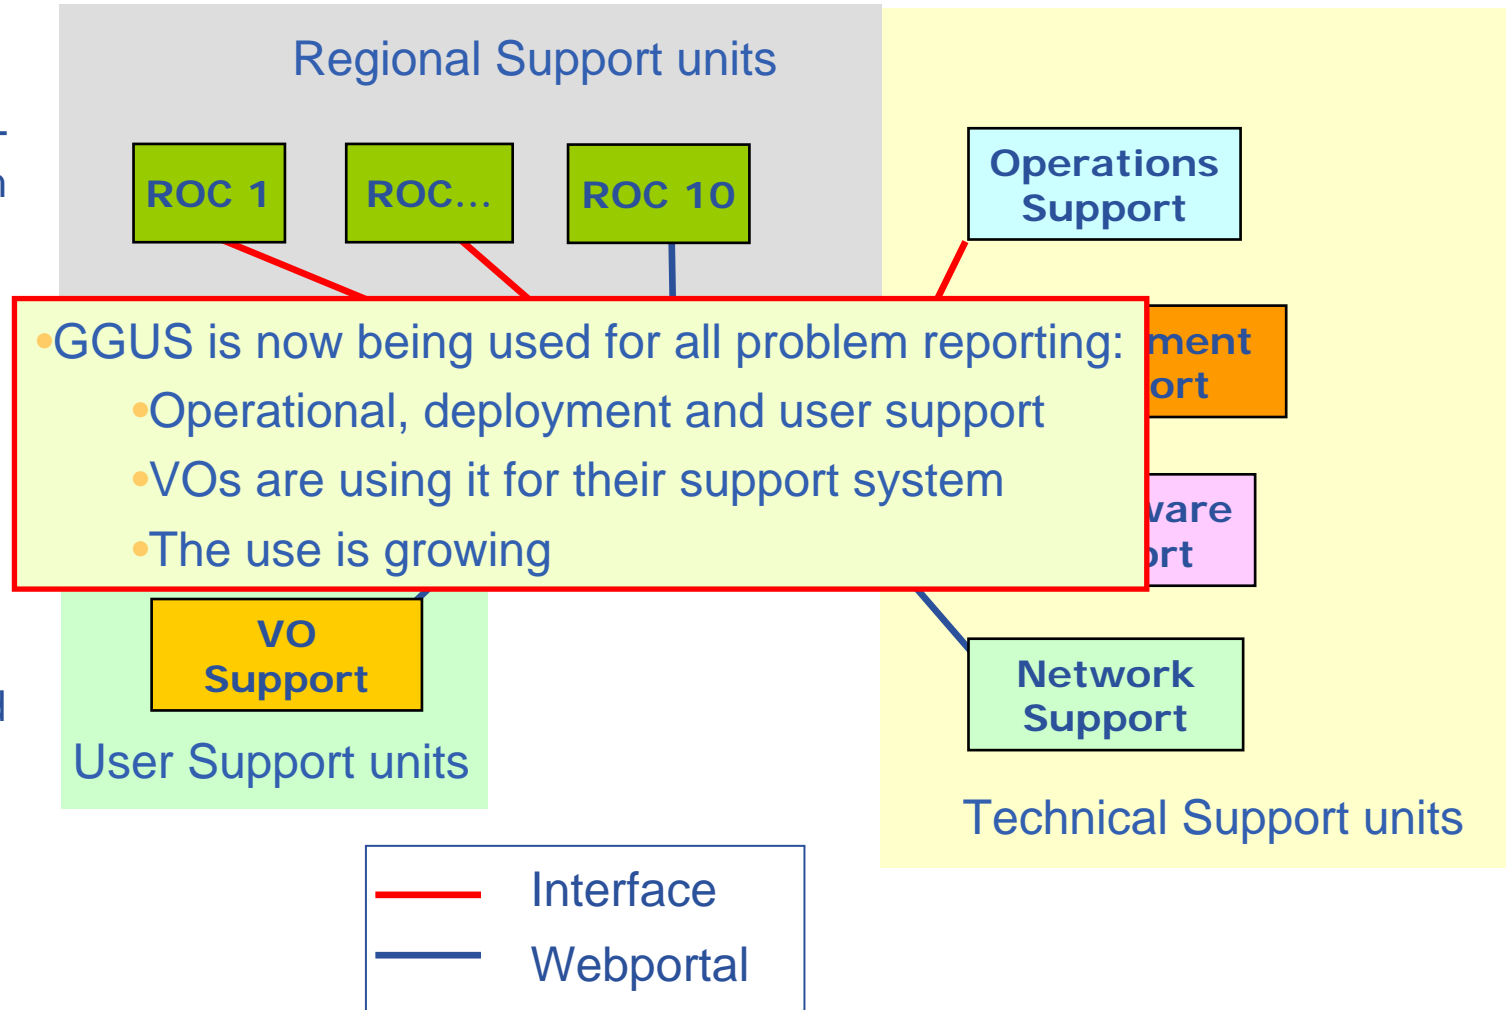
- Based on LCG-2.7.0, and
- gLite workload management
- Other gLite components (not in the distribution but provided as services):
 - AMGA, Hydra, Fireman
 - gLite-IO





"Regional Support with Central Coordination"

The ROCs, VOs and other project-wide groups such as the middleware groups ([JRA](#)), network groups ([NA](#)), service groups (SA) are connected via a central integration platform provided by GGUS.



- **Scale of the infrastructure**
 - Has grown steadily during the project
 - Now slowed – expansion with related projects
- **Sustained real production use of the infrastructure**
 - Is supported by the operations teams
- **Maturing but evolving operations procedures**
 - Dealing with all aspects of operations
- **User support**
 - GGUS is becoming the central coordination point, use is growing
- **Middleware distribution**
 - Now clear how to evolve the production service
 - Convergence between existing LCG-2.x and gLite-1.x
- **Progress on interoperability and interoperation**
 - With OSG significant progress,
 - Progress with ARC

- **LHC VOs must achieve reliable production and analysis in 2006**
 - Will be making significant use of resources
- **Consolidate and improve existing services: Focus on**
 - Reliability, robustness
 - Manageability
 - Performance, scalability
 - Evolution or replacement of services driven by needs of application (or security/manageability)
- **Expand grid operations**
 - Spread expertise to ROCs
 - Collaboration with OSG, A-P, etc. and related projects
 - Start to negotiate SLAs
- **New applications**
 - Must bring resources – show commitment
 - Resource sharing and negotiation – must become streamlined
 - Will need a mechanism for cost/credit for use of resources

