SA1 Status Report

EGEE Grid Operations & Management

Ian Bird

SA1 Activity Leader

IT Department, CERN

1st EU Review of EGEE-II CERN *15-16th May 2007*

www.eu-egee.org

Information Society and Media

SA1 in Numbers

SA1 Partners

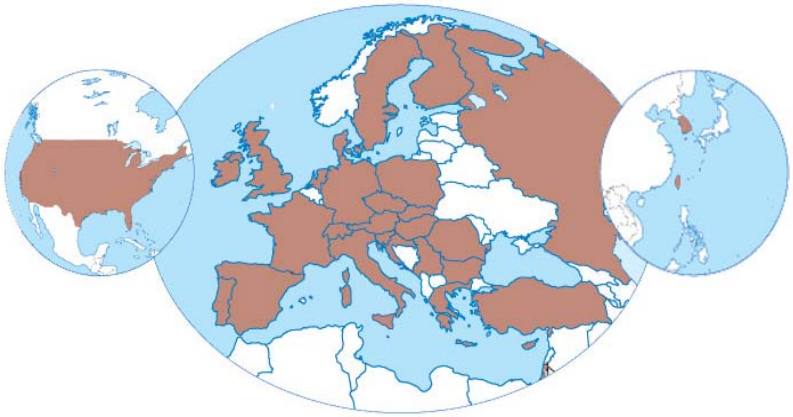

EGEE-II Budget

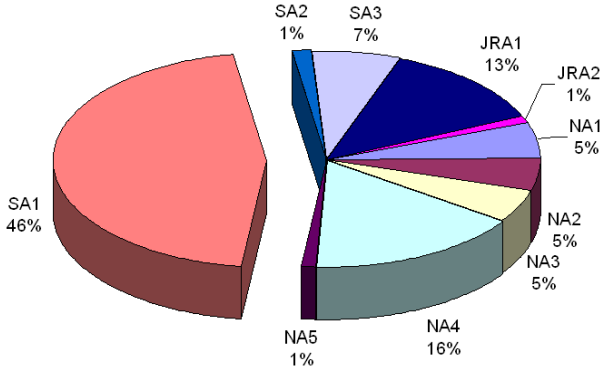

Manpower: 61 partners, 29 countries, 228 FTE

Short Name	Country	Total (PMs)
CERN	Switzerland	480
JKU	Austria	24
UIBK	Austria	25
CESNET	Czech Rep.	72
CSC	Finland	12
KFKI-RMKI	Hungary	68
NIIF	Hungary	28
CYFRONET	Poland	108
ICM UW	Poland	36
PSNC	Poland	24
IISAS	Slovakia	36
JSI	Slovenia	36
TCD	Ireland	48
CCLRC	UK	210
UEDIN	UK	72
Imperial	UK	24
CSCS	Switzerland	24
CEA	France	24
CGG	France	24
CNRS	France	504
CS SI	France	48
DESY	Germany	72
FhG/SCAI	Germany	60
FZK	Germany	144
GSI	Germany	60
INFN	Italy	708
UKBH	Denmark	24
FOM	Netherlands	96
SARA	Netherlands	96
VR	Sweden	132
IHEP	Russia	78
IMPB RAS	Russia	48
ITEP	Russia	78
JINR	Russia	90
KIAM RAS	Russia	42
PNPI	Russia	36
RRC KI	Russia	60
SINP MSU	Russia	96
IPP-BAS	Bulgaria	96
UCY	Cyprus	96
GRNET	Greece	216
TAU	Israel	120
ICI	Romania	120
LIP	Portugal	96
CSIC	Spain	84
PIC	Spain	120
TID	Spain	48
ENEA	Italy	21
UNICAL	Italy	20
UNILE	Italy	20
UNINA	Italy	20
RED.ES	Spain	48
CESGA	Spain	84
IPB	Serbia and Montenegro	48
TUBITAK-ULAK	Turkey	48
RUG	Netherlands	12
Glasgow	UK	48
UNIMAN	UK	48
Oxford	UK	48
ASGC	Taipei	135
Srce	Croatia	32
Total		**5475**

The EGEE Infrastructure

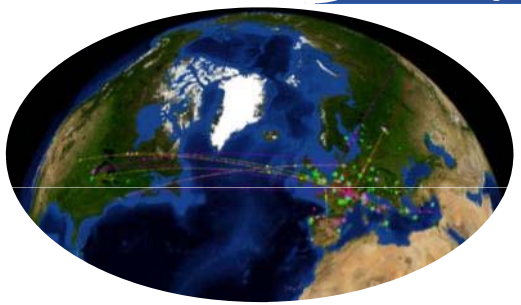

Test-beds & Services

- Production Service
- Pre-production service
- Certification test-beds (SA3)
- Training infrastructure (NA4)

Support Structures & Processes

- Operations Coordination Centre
- Regional Operations Centres
- Global Grid User Support
- EGEE Network Operations Centre (SA2)
- Operational Security Coordination Team
- Training activities (NA3)

Security & Policy Groups

- Joint Security Policy Group
- EuGridPMA (& IGTF)
- Grid Security Vulnerability Group
- Operations Advisory Group (+NA4)

CPU, Sites, countries

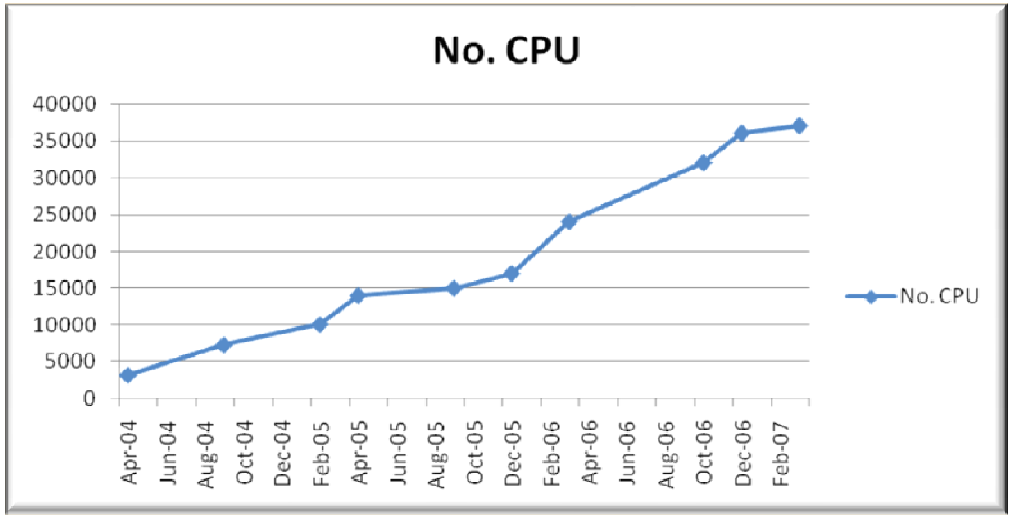

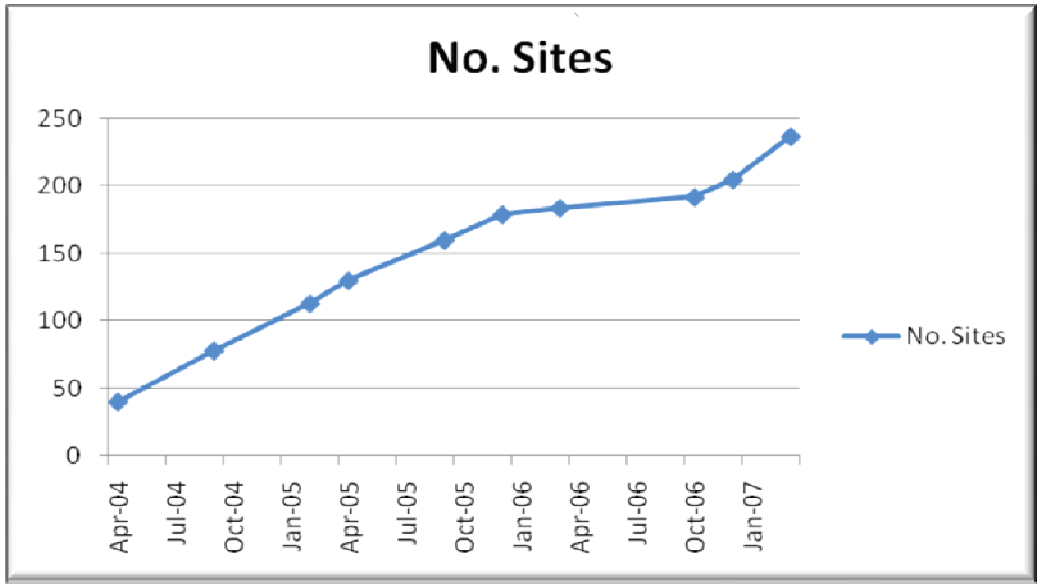

ROC	Partner - DoW	Partner - actual	Total	% non partner
CERN	1800	3548	5943	40%
France	1252	2550	2700	6%
De/CH	1852	2695	3364	20%
Italy	2280	3539	3628	2%
UK/I	2010	4527	7720	41%
CE	1163	1622	1875	13%
NE	1860	2473	3031	18%
SEE	1289	2552	2568	1%
SWE	898	1535	1593	4%
Russia	445	527	583	10%
A-P	801	841	1632	48%
Total	**15650**	**26409**	**34637**	**24%**

CPU, countries, sites

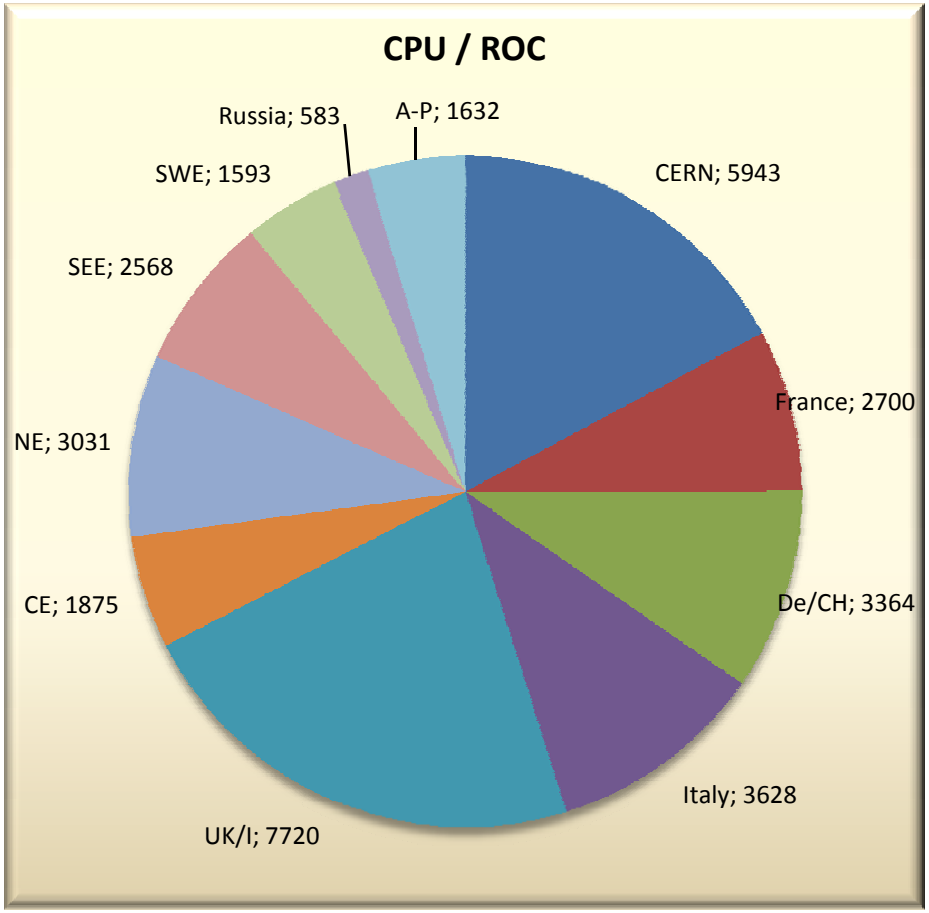

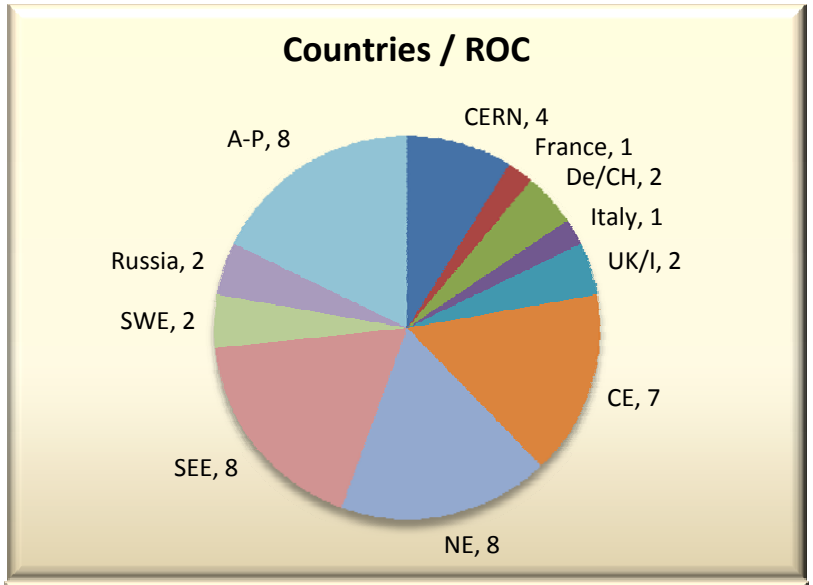

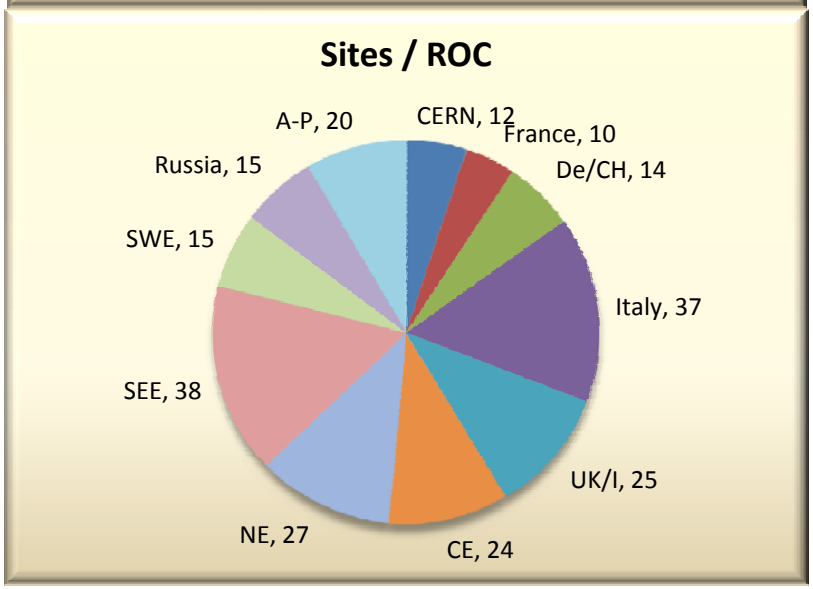

- 35000 CPU
- 45 countries (31 partner countries)
- 237 sites (131 partner sites)

Workload

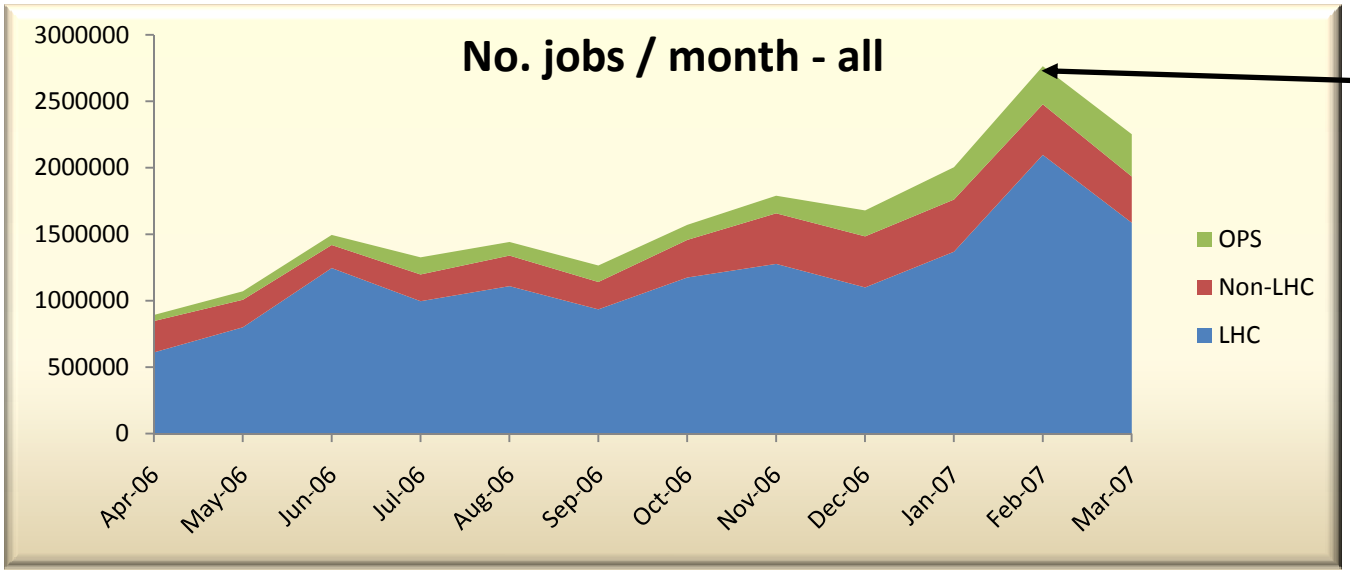

98000 jobs/day

13000 jobs/day

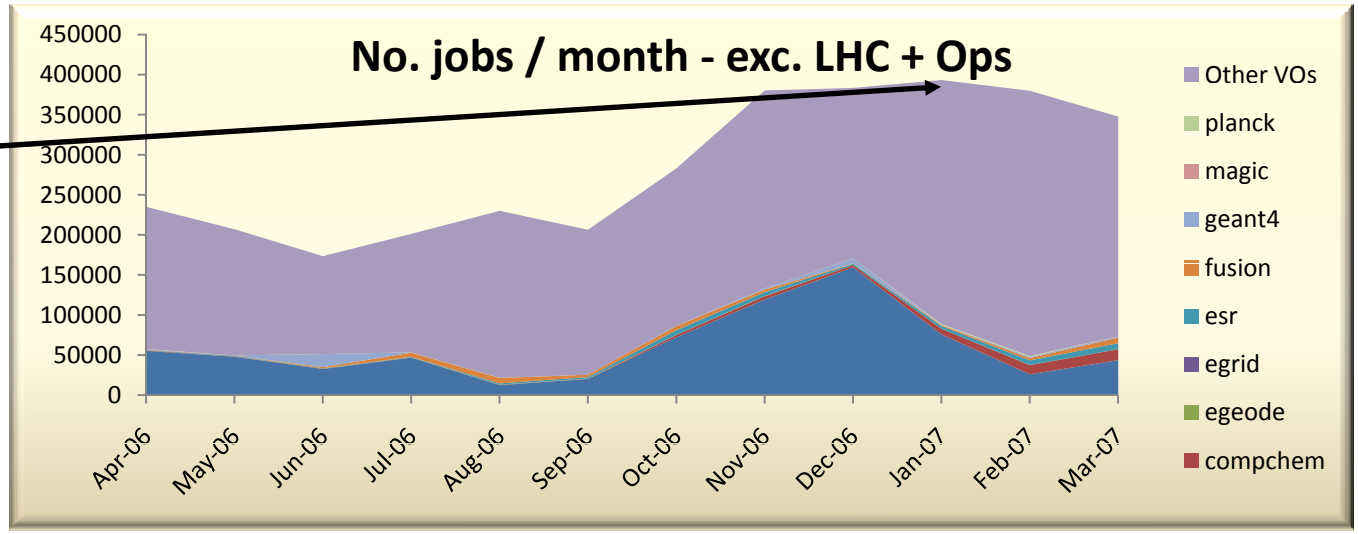

CPU time delivered

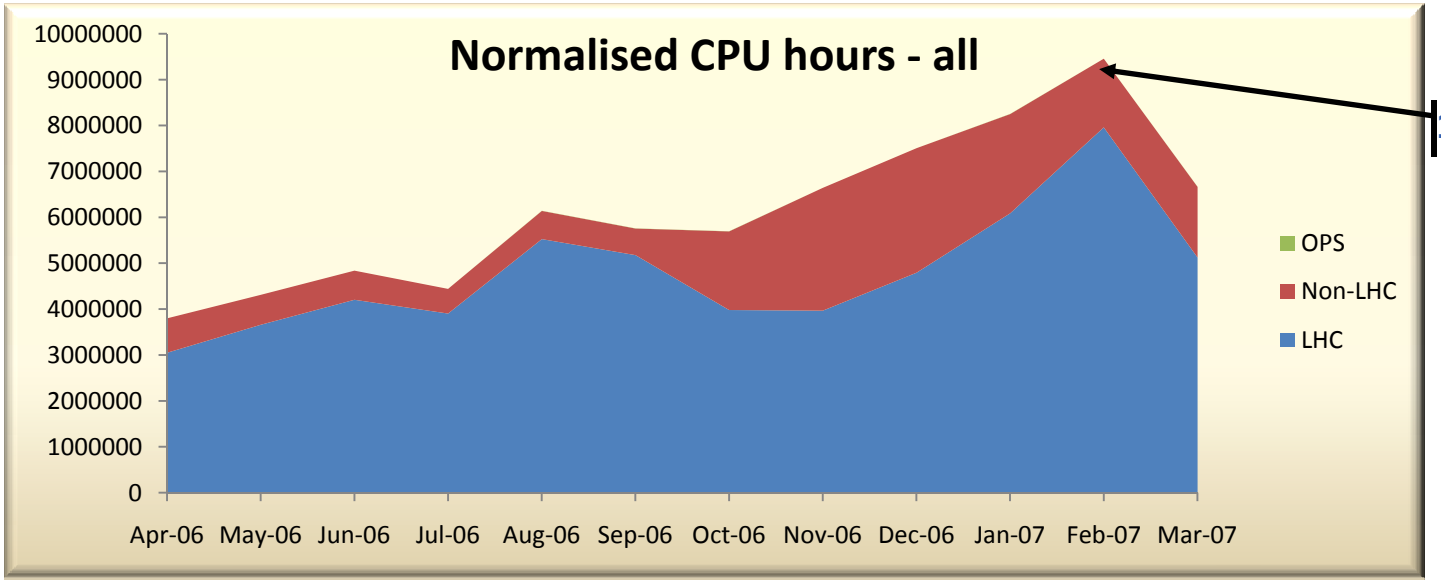

14000 CPU-month/month

3600 CPU-month
~ 1/3 of total

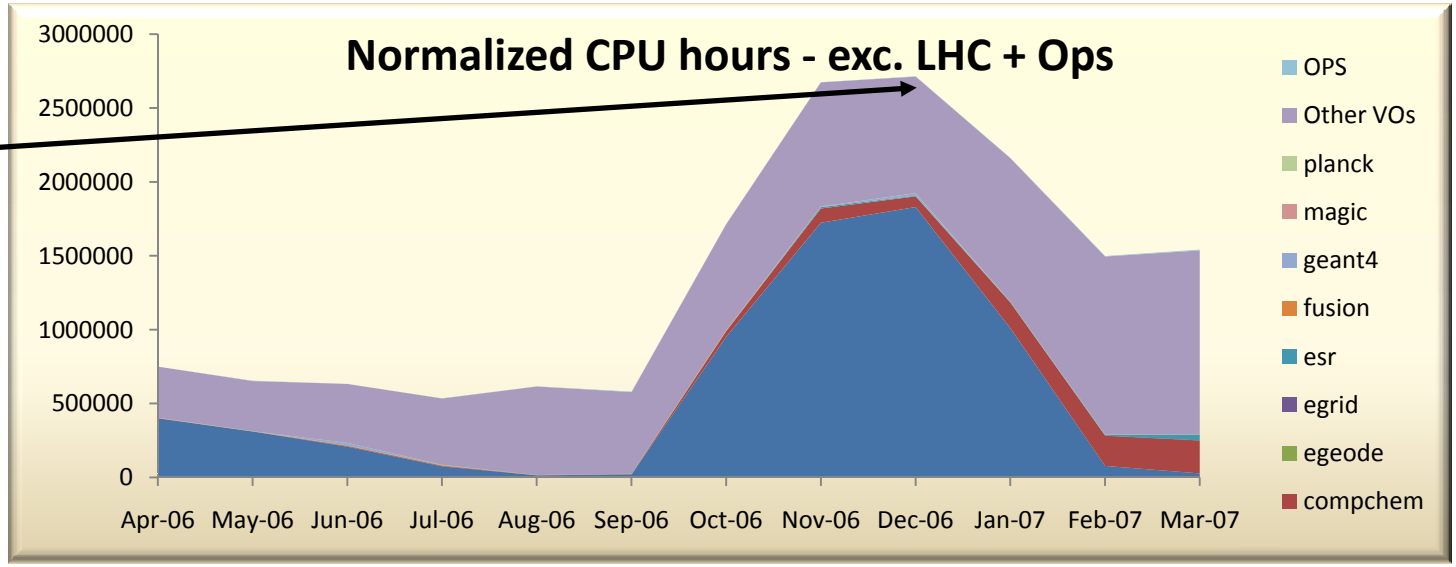

Overall load

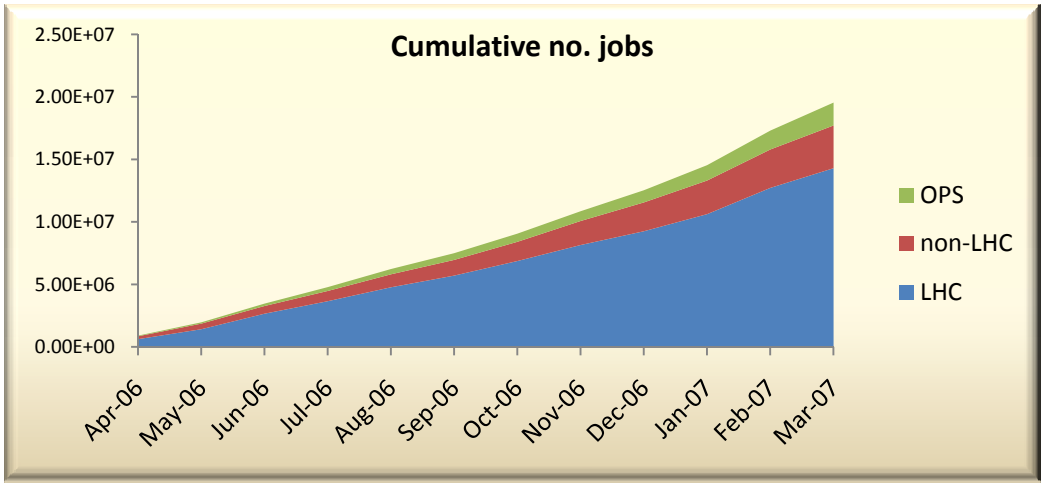

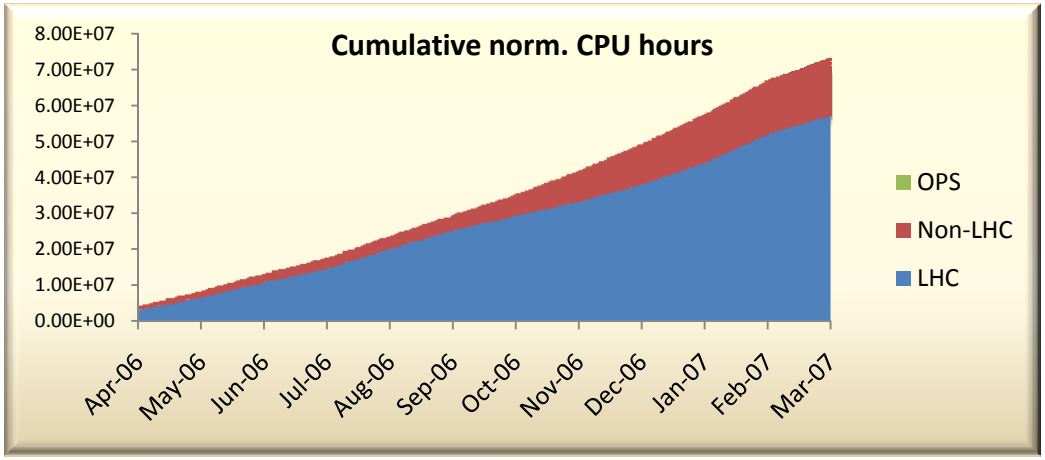

- **19.6 million jobs run in 1st year of EGEE-II**
  - 56000 per day sustained average
  - Peak of 98000
  - Non-LHC 13500 /day
    - Level of total in EGEE in 2005
- **8400 CPU-years delivered in 1 year**
  - ~1/3 of total available sustained over the year
  - Peak of 50% of available in Feb ’07
  - ~1/3 of total was non-LHC in Dec ‘06

Regional distribution

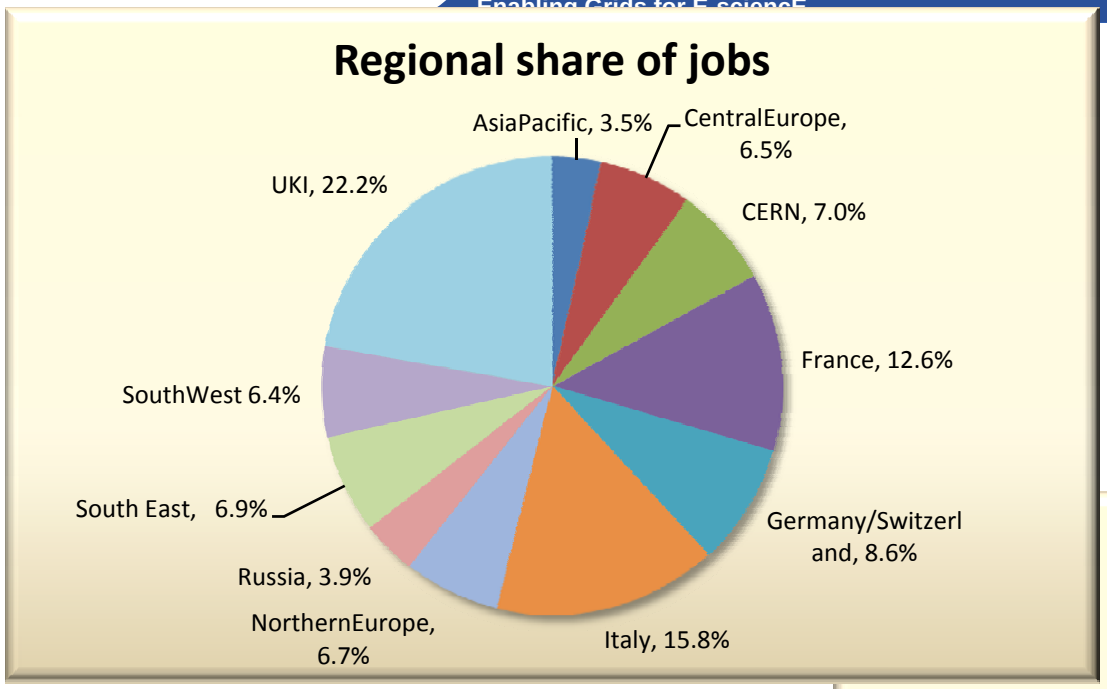

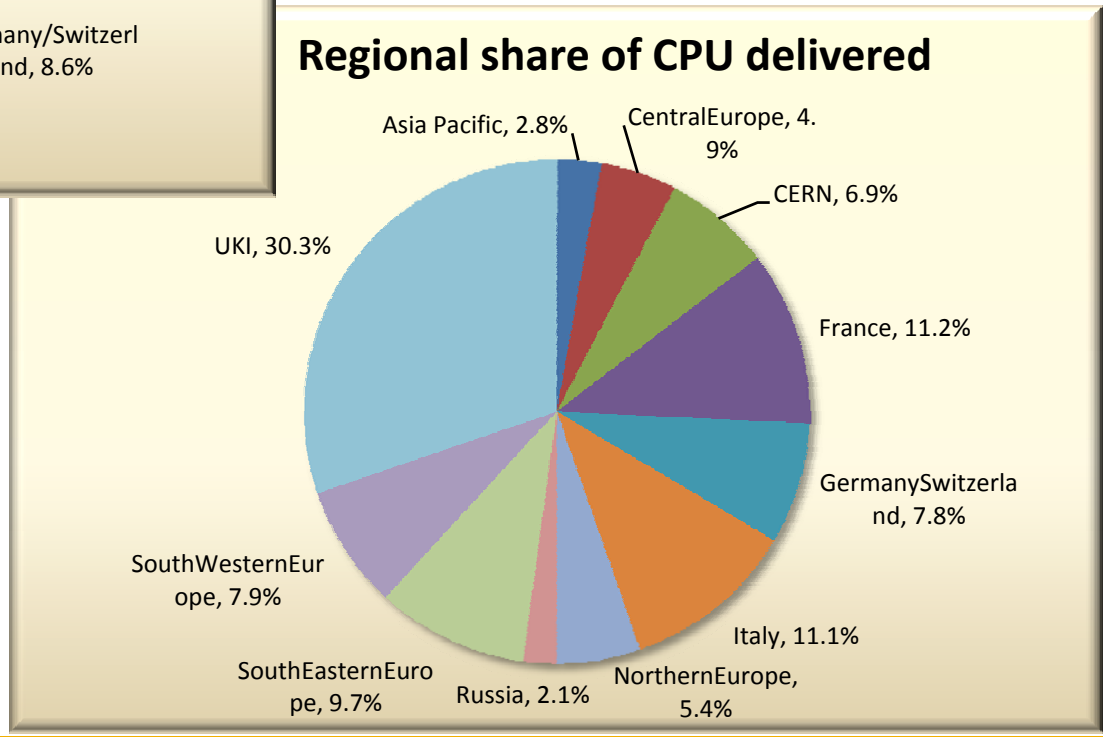

Use of the infrastructure

Enabling Grids for E-sciencE

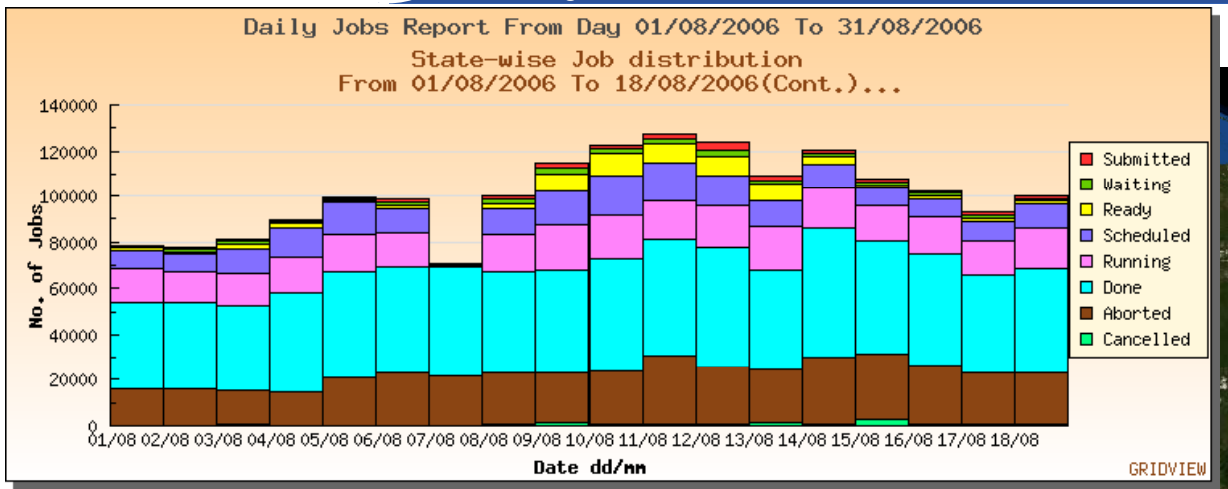

Scheduled = 17356
Running = 18359

>20k jobs running simultaneously

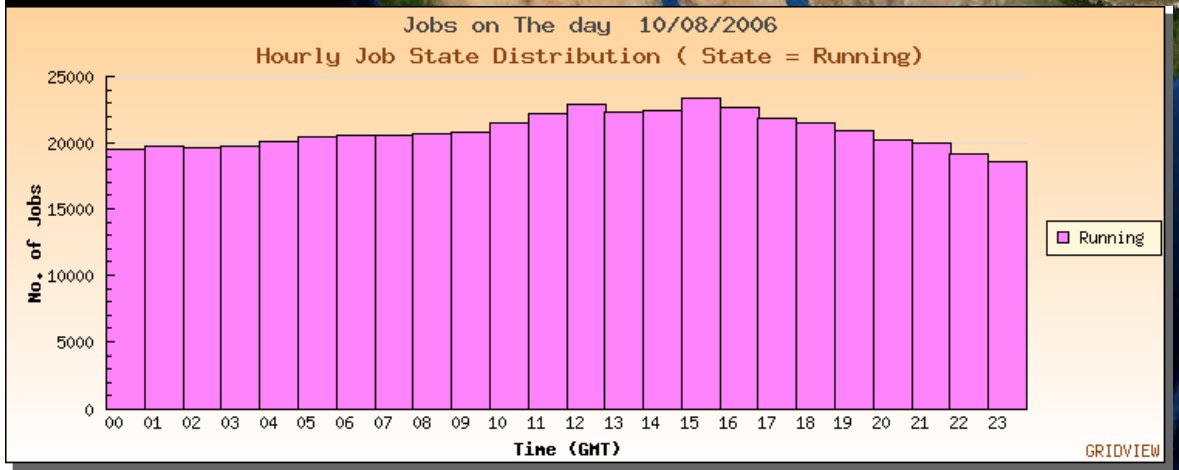

08:54:05 UTC

GridPP
UK Computing for Particle Physics

Use for massive data transfer

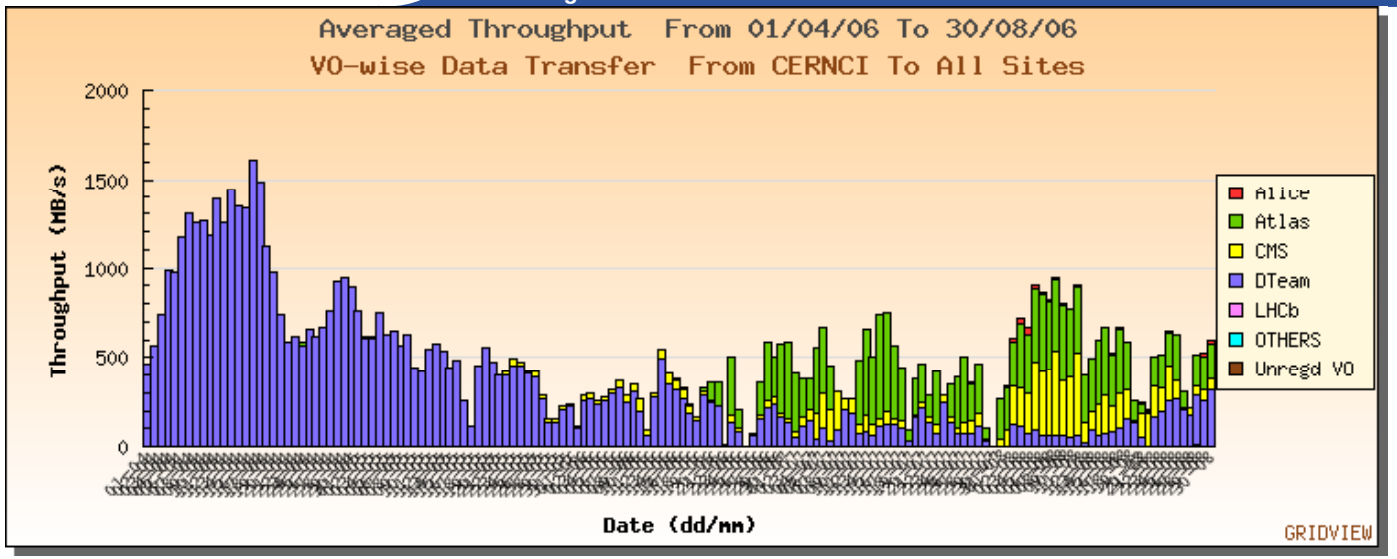

Large LHC experiments have demonstrated transfers ~ 1PB/month each
Reach ~1GB/s aggregate data rates with real workloads

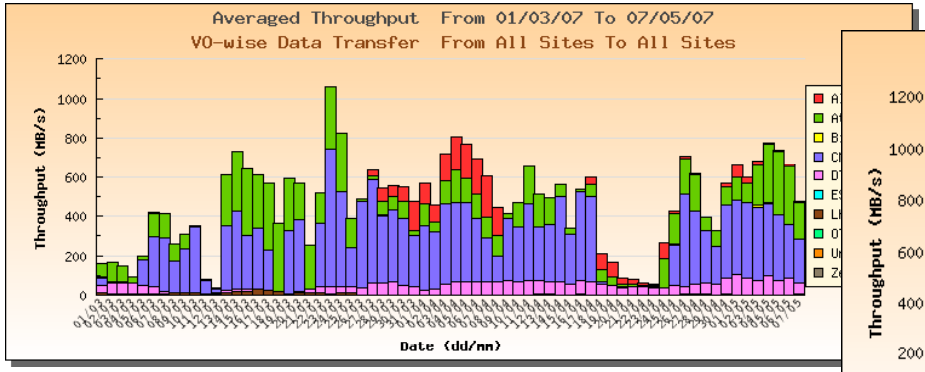

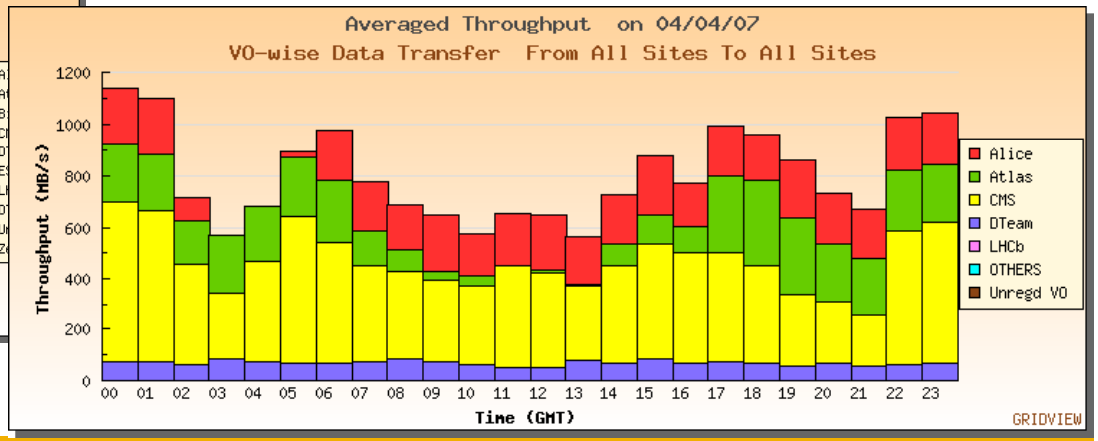

Pre-production service

- **Pre-production service is now ~ 27 sites in 16 countries**
- **Provides access to some 3000 CPU**
  - Some sites allow access to their full production batch systems for scale tests
- **Sites install and test different configurations and sets of services**
  - Weekly update cycle
  - Try to get good feeling for the quality of the release or updates before general release to production
  - Larger sites gain experience on PPS before going to production.
- **Services may be initially demonstrated in this environment**
  - Before further development
- **New VO-s: adapt their applications & gain experience**
  - (e.g. DILIGENT)

Grid Operations

- **Grid Operator on Duty ("CoD")**
  - Teams from 10 of 11 ROCs participate
  - 5-weekly rotations: each week 1 team primary and 1 team backup
  - Critical activity in maintaining usability and stability of sites
  - Important tools
    - Site Availability Tests (SAM)
    - Information system monitoring
    - GGUS system for trouble ticket management
  - Portal for operations : https://cic.gridops.org

- **Significant work on operations procedures**
  - Evolved throughout EGEE and EGEE-II
  - Contribute to establishment of regional grid infrastructures through related projects – well beyond Europe now

User Support

- **GGUS – now well established**
  - Used more and more
  - 10 of 11 ROCs provide dedicated effort to manage the process – similar to operator on duty teams
  - Development plan (DSA1.1) and assessment of progress (MSA1.8) delivered
  - Set up user support advisory groups to steer the priorities

- **GGUS tool used for all support activities**
  - Interlinks many local ticketing systems

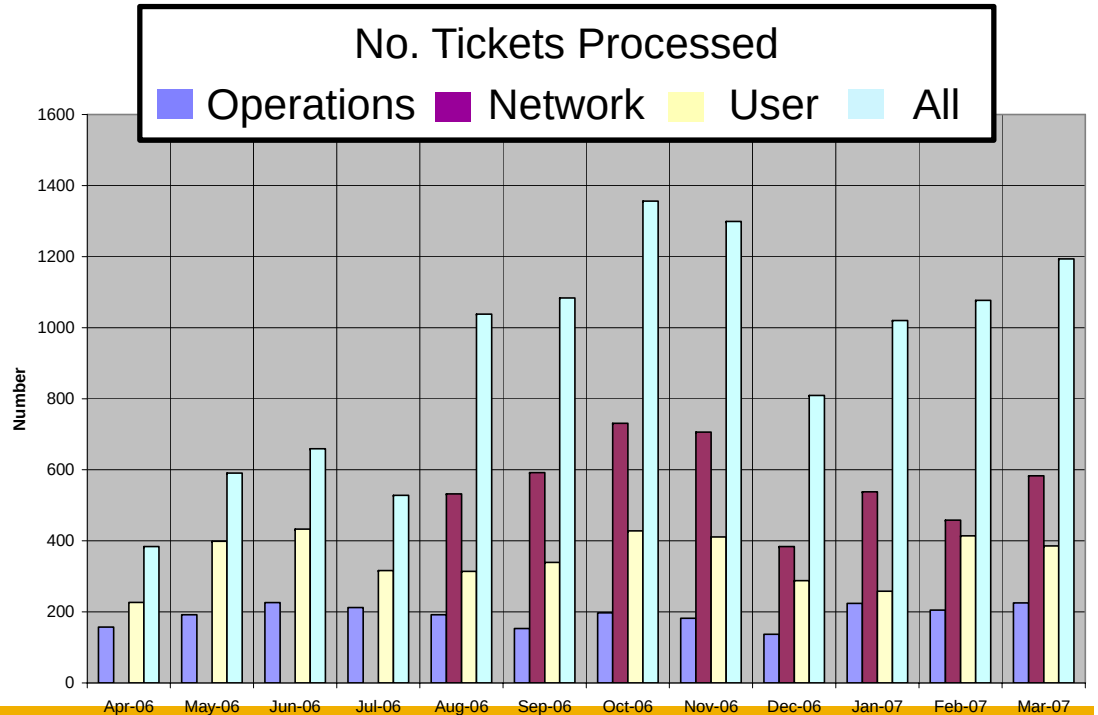

Joint Security Policy Group

- **Joint Security Policy Group (JSPG)**
  - Produces and maintains security policy and procedures
    - for EGEE, OSG, NDGF, WLCG, and other EU Grid infrastructures
- **Main successes**
  - Achieved common policy between EGEE and OSG (for interoperation)
  - New Grid Site Operations Policy (MSA1.3)
  - Updated top-level Security Policy (MSA1.7)
  - Grid User AUP accepted by eIRG as good approach
- **Current work**
  - New policy addressing User-level Accounting (data privacy issues)
  - New policy on VO and Grid service responsibilities
- **Issues**
  - Much of the effort is based on volunteers (progress can be slow)
  - Still need to make policies more general and simpler
  - Need to involve more NGI's (for better interoperation)

Operational Security Coordination Team

- **Operational Security Coordination Team (OSCT) focuses on:**
  - Incident Response & improvement
  - Security Monitoring
  - Best practice for system managers
    - Anticipate using results of ISSeG project
  - Pan-regional security coordination
- **Main improvements**
  - Team members are on rota to ensure timely response to operational issues
  - Security Service Challenges remain very useful
  - ROCs are more involved in pan-regional activities
  - The team is now more cohesive
- **Issues**
  - Most ROCs are unable to deliver agreed efforts
  - Several ROCs are unable to contribute to any activity other than the rota
  - Communication with peer Grids and NREN CERTs has improved but remains weak

The Grid Security Vulnerability Group

- **Aim: "to incrementally make the Grid more secure and thus provide better availability and sustainability of the deployed infrastructure"**
- **Process and Infrastructure in place – DSA1.3**
- **Key: Risk Assessment Team performs a risk assessment on each issue according to an agreed (objective) strategy**
- **A Target Date (TD) for resolution is set according to the risk**
- **Each issue is placed in one of 4 risk categories**
  - Extremely Critical (TD = 2 working days)
  - High (TD = 3 weeks)
  - Moderate (TD = 3 months)
  - Low (TD = 6 months)
- **An advisory is issued:**
  - with the software patch, or
  - on the TD if a patch is not available
- **102 potential issues in total have been submitted and handled by the group**
  - 49 closed (20 fixed, 16 invalid, 2 general concerns we consider to be adequately addressed, 6 operational/OSCT informed, 4 duplicate, 1 documentation changes)
  - 53 open (37 awaiting bug fix, 16 more general concerns in discussion)

Interoperability/interoperation

- **Well established with OSG**
  - In production use by CMS – submits work to OSG from EGEE
- **Weekly operations meetings attended by OSG staff**
- **Operations workshops have sessions on interoperability**
- **Processes set up with OSG for operations and user support workflows**
- **OPS VO defined to support joint operations – for testing/monitoring use**
- **OSG use SAM concept and will report site tests into SAM (for LCG)**
- **Agree to set up an interoperability test-bed as part of both projects' certification activities – seen as essential now**
- **Interoperation also with NDGF for LCG – through EGEE mechanisms**
- **Contacts with NAREGI established – expect to increase**

Worldwide Grid Infrastructures

Enabling Grids for E-sciencE

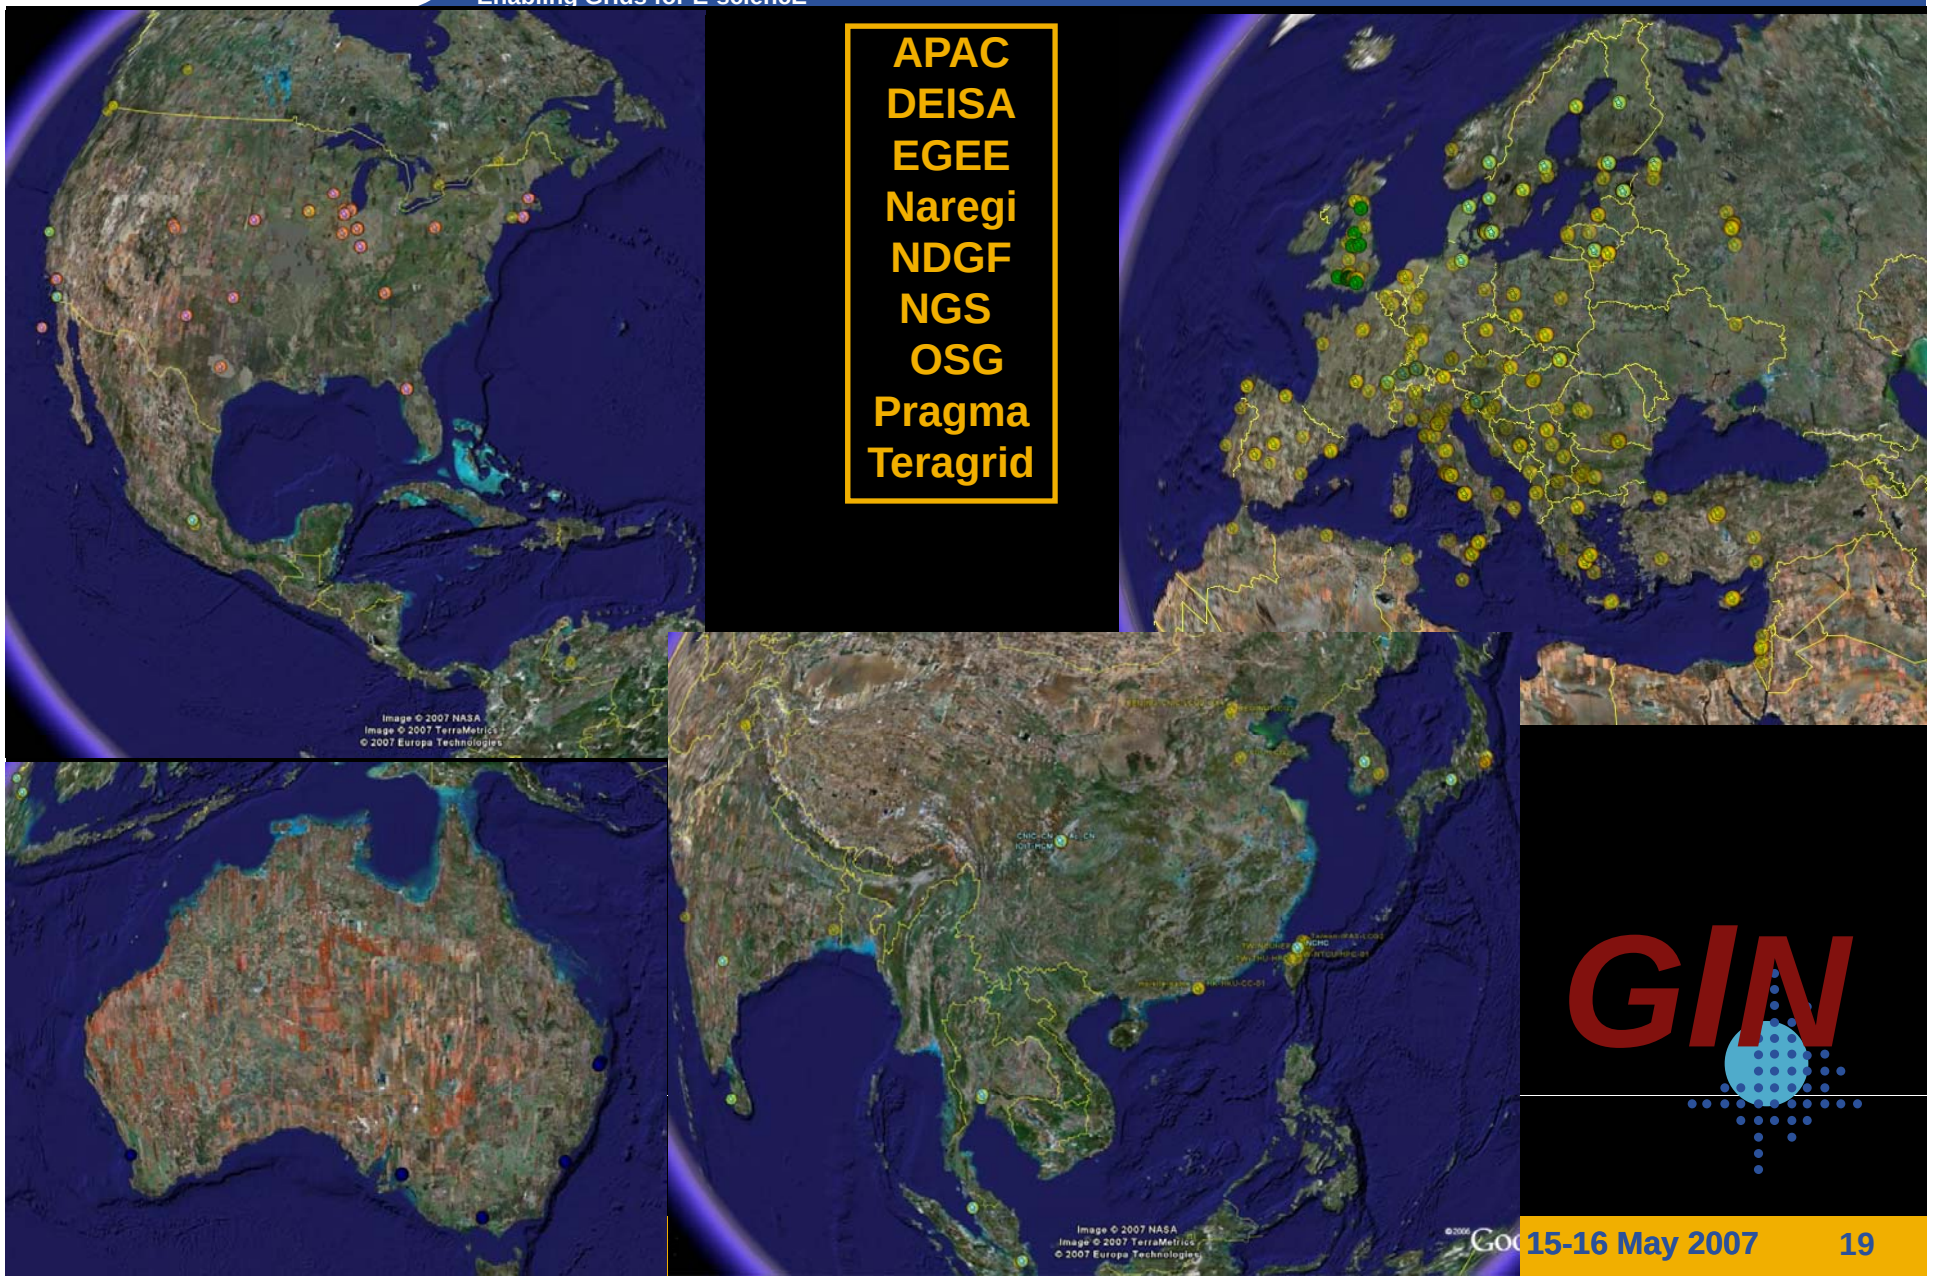

Monitoring

- **3 groups set up**

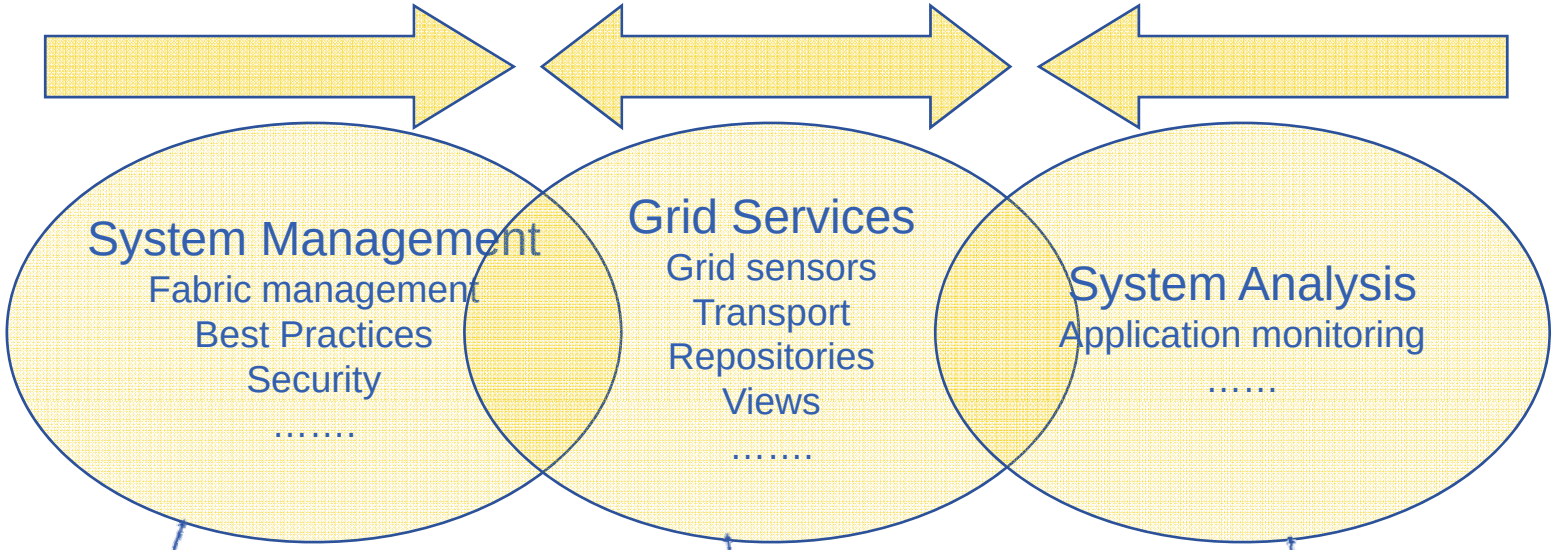

- " … improving system management practices,
- Provide site manager input to requirements on grid monitoring and management tools
- Propose existing tools to the grid monitoring working group
- Produce a Grid Site Fabric Management cook-book
- Identify training needs

- "… To help improve the reliability of the grid infrastructure …"
- " … provide stakeholders with views of the infrastructure allowing them to understand the current and historical status of the service …"

- " … to gain understanding of application failures in the grid environment and to provide an application view of the state of the infrastructure …"

Monitoring Data Flow

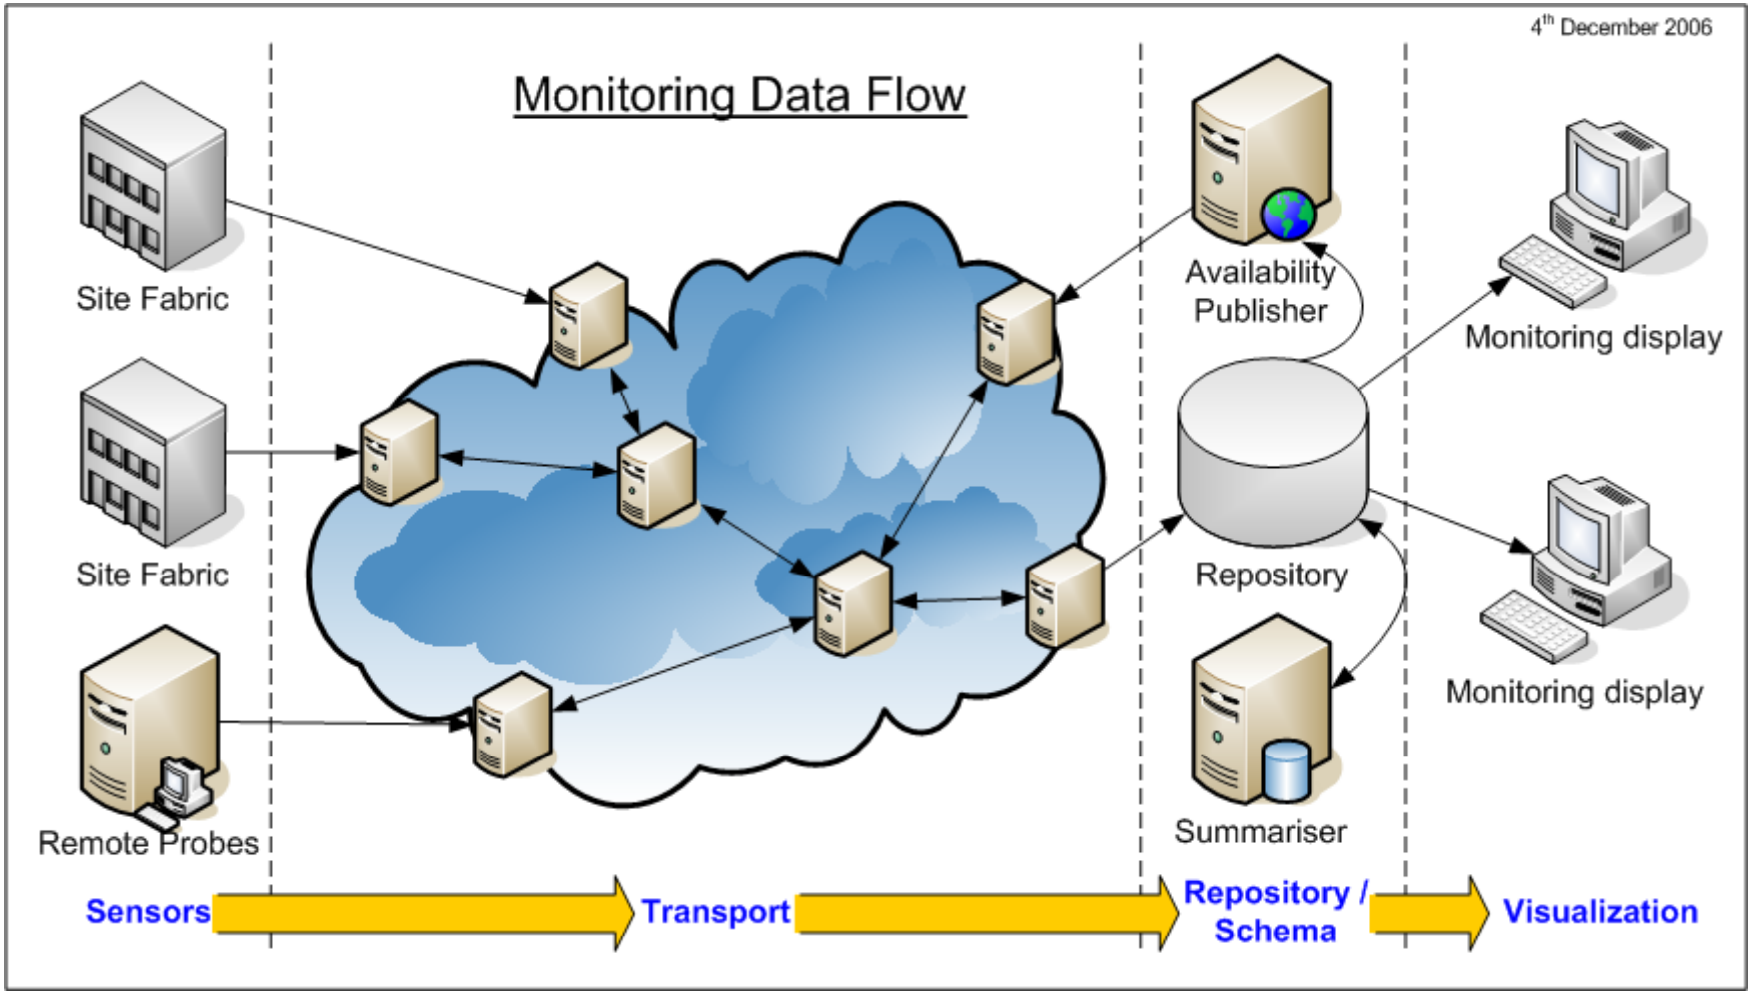

Site Metrics Publication

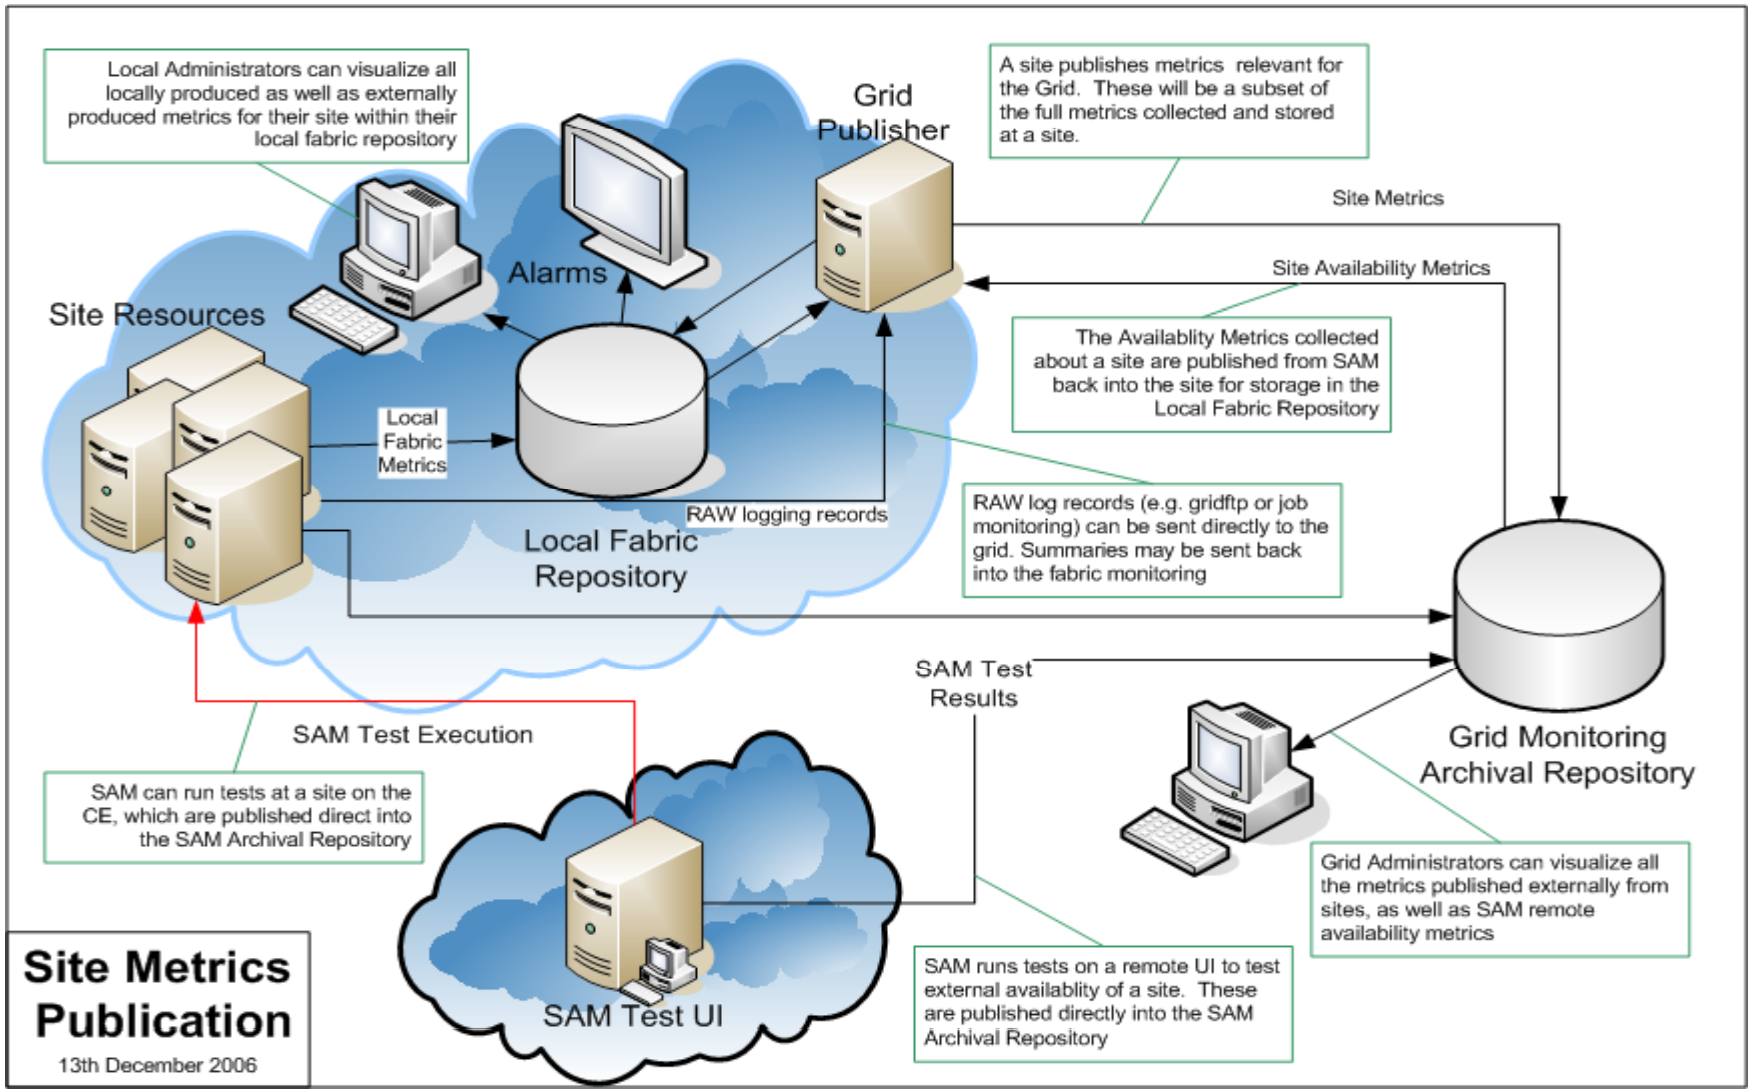

Experiment Dashboard

INPUT
Multiple sources of information
Increasing the reliability
Providing both global and very detailed view

OUTPUT
Providing output in various formats
(Web pages, xml, csv, image formats)

Collect data of VO interest coming from various sources
Store it in a single location
Provide UI following
Define alarm conditions

Can satisfy users with various roles:
Generic
on the Grid
Site administrator
VO manager, production or analysis group coordinator, data transfer coordinator…

This will be shown in the demo session

Can be used by various clients
applications

report monitoring information

Potentially other Clients:
PANDA, ATLAS production
<XML,CSV, image formats>

Job monitoring for CSA06

Enabling Grids for E-sciencE

any user
any site
any ce
any submissiontool
any dataset
any application
any rb
any activity
any grid
☐ unk ☐ pend ☐ run ☐ term
☐ done ☐ canc ☐ abort ☐ g-unk
☐ succ ☐ fail ☐ a-unk
☐ donesuccess
2006-11-08 19:44:55
to
2006-11-09 19:44:55
sort by activity
bars in the plot
submit

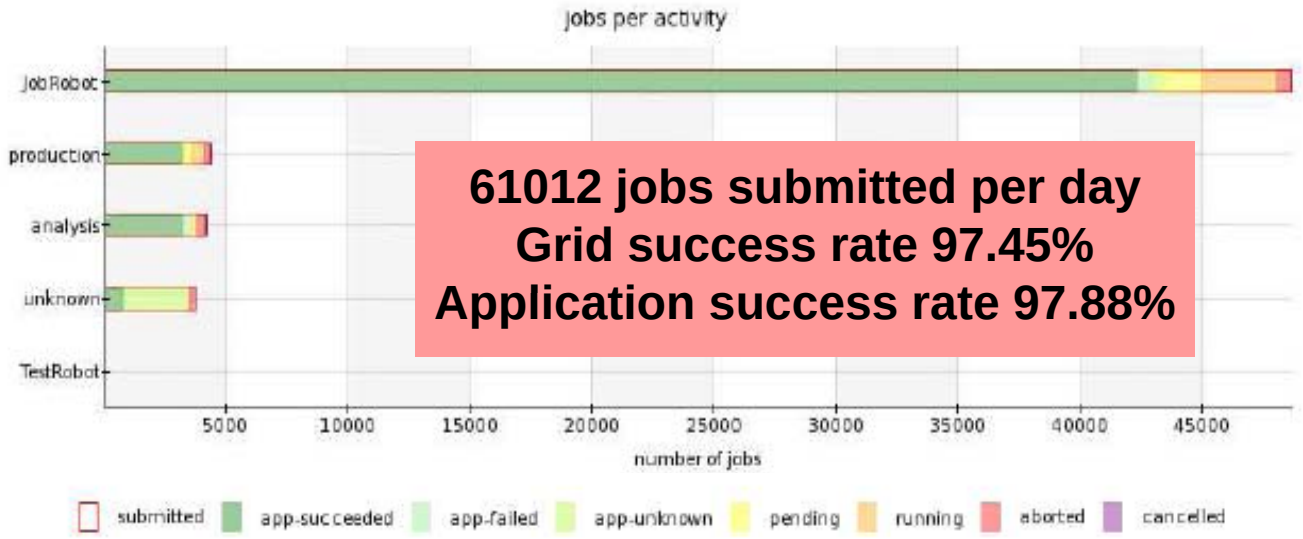

61012 jobs submitted per day
Grid success rate 97.45%
Application success rate 97.88%

	dashboard					grid					application				overall	
activity↓	Sub	Unk	Pend	Run	Term	Done	Canc	Abort	Unk	Grid%	Succ	Fail	Unk	App%	D/S	Overa
JobRobot	486980		1683	3100	439154	084658		539	2472	98.7	42320547		1048	98.72	401039	91.32
TestRobot	12	0	0	0	12	0	0	0	12	0	0	12	0	0	0	0
analysis	4206	0	182	21	4003	2592	113	291	1007	89.91	3181	406	416	88.68	2178	54.41
production	4351	0	299	543	3509	1016	68	184	2241	84.67	3190	5	314	99.84	976	27.81
unknown	3745	0	9	16	3720	3425	21	238	36	93.5	767	48	2905	94.11	655	17.61
total	*610120*		*2173*	*3680*	*551594*	*787926*	*0*	*1252*	*5768*	*97.45*	*49458101*	*8*	*4683*	*97.98*	*439127*	*9.61*

query took 0.78 seconds.

Note: How job status and success rate are calculated?

Accounting

- **Accounting system set up by UK/I – now well established – all sites reporting into it**
  - Now starting to deploy a version that reports by user
  - User DN is encrypted for privacy
  - Policy (in draft) that defines who can access what information and for what purpose

- **Storage accounting – prototype available now**
  - Schema has been defined
  - Uses information system to publish available and used storage space data, for different classes of storage
  - Sensor queries the BDII and stores into R-GMA and the APEL system
  - Portal to query the data is based on the CPU accounting portal

VO Resource Manager view

- Table shows CPU, WCT and Job Eff. of the Top 10 Anonymised Users
- Breakdown of Usage: DN / VO / Group / Role

EGEE View | VO MANAGER View | VO MEMBER View | SITE ADMIN View | USER View

January 2006 - December 2006.

The following table shows the Usage of the Top 10 Users ordered by Normalised CPU time and the Total Usage of the Other Users. A detailed view can be obtained by selecting an individual user.

Top 10 Users ordered by Normalised CPU time

	User	Jobs		CPU time		*Norm. CPU time*		WCT		Norm. WCT		CPU Efficiency	Avg. CPU time
#	ID	#	%	Hrs	%	*Hrs*	*%*	Hrs	%	Hrs	%	%	Hrs
1	2daf892f6a1f68d0	15,629	25.2%	120	24.3%	*88*	*22.6%*	1,314	16.5%	990	16.0%	9.1	0.01
2	671e5b277fa6d1ac	20	0.0%	98	19.8%	*87*	*22.3%*	103	1.3%	91	1.5%	95.1	4.90
3	4cd78d6b029f7050	7,773	12.5%	80	16.2%	*55*	*14.1%*	924	11.6%	765	12.4%	8.7	0.01
4	5d61bd2201bec5f9	7,816	12.6%	52	10.5%	*41*	*10.5%*	331	4.2%	250	4.0%	15.7	0.01
5	2a153e141e98f06a	1,950	3.1%	36	7.3%	*36*	*9.2%*	45	0.6%	46	0.7%	80.0	0.02
6	62d3866c3c8260d6	39	0.1%	27	5.5%	*16*	*4.1%*	31	0.4%	18	0.3%	87.1	0.69
7	2756dfcb65975a47	95	0.2%	19	3.8%	*15*	*3.8%*	22	0.3%	18	0.3%	86.4	0.20
8	748206ea352cce31	467	0.8%	12	2.4%	*14*	*3.6%*	15	0.2%	16	0.3%	80.0	0.03
9	2d04c1be5d64c1b8	3	0.0%	11	2.2%	*13*	*3.3%*	26	0.3%	21	0.3%	42.3	3.67
10	2d72edb26620a697	83	0.1%	9	1.8%	*6*	*1.5%*	125	1.6%	79	1.3%	7.2	0.11
	Others (DN known)	25,073	40.5%	18	3.6%	*10*	*2.6%*	4,695	58.9%	3,527	57.1%	0.4	0.00
	Others (DN unknown)	3,021	4.9%	12	2.4%	*9*	*2.3%*	343	4.3%	356	5.8%	3.5	0.00
	Total	**61,969**		**494**		***390***		**7,974**		**6,177**		**6.2**	**0.01**

Click here for a csv dump of this table

Key: 0% <= eff < 75%; 75% <= eff < 90%; 90% <= eff < 100%; eff >= 100% (parallel jobs)

Storage Accounting Display

- **Storage units are 1TB = 10^6 MB**

Disk Used - Italy

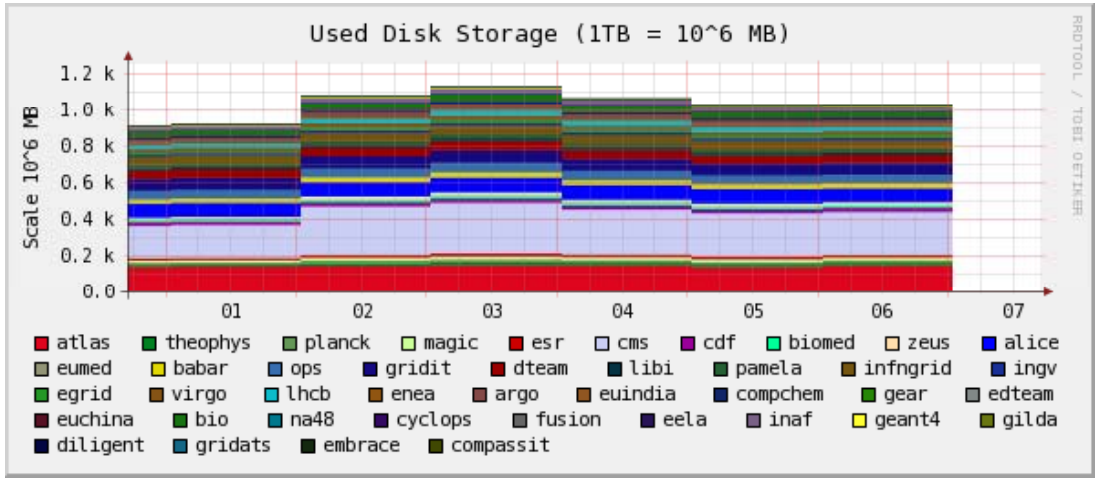

Tape Used – UK/I

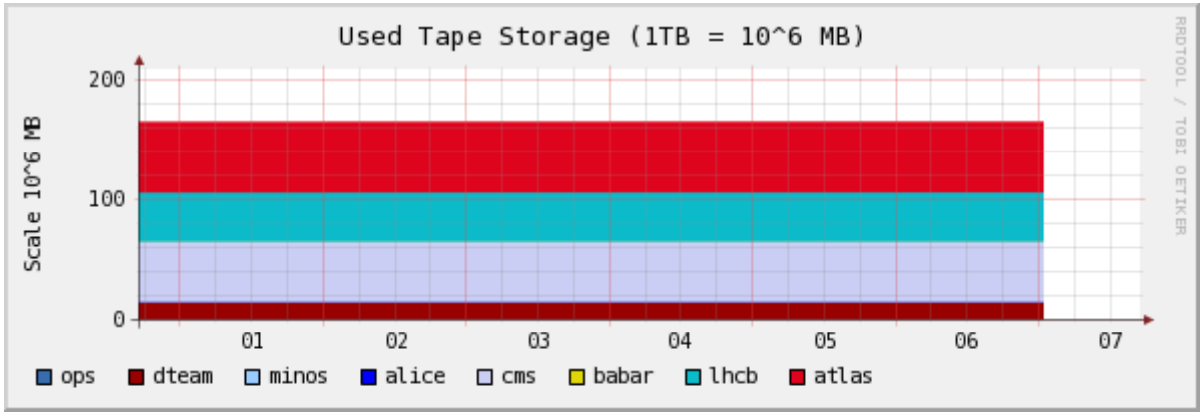

Network Performance Monitoring

- **NPM is now an SA1 task**
  - With its own GGUS Support Unit
  - ENOC the main user, contributing to requirements
    - Interest from ROCs during recent network outage
- **e2emonit, our end-to-end NPM framework, was:**
  - Certified and included in gLite 3.0
  - Deployed on 6 sites of the PPS
  - Selected by BalticGrid for deployment
- **Collaboration with GÉANT2 peers (JRA1) continues:**
  - NPM brokers into EGEE the GÉANT2 PerfSonar NPM data
  - We consciously use the same NMWG schema
    - …and influence OGF-NMWG in tandem
  - We exchange experiences and strategy plans
- **Dissemination:**
  - http://www.egee-npm.org/
  - 2 conference papers and 1 EGEE TR

Partner reviews

- **SA1 goals in doing reviews**
  - Alternative to partner metrics – Regions are different – no simple measure will work – need to look in sufficient detail to understand what is going on
  - Do a review of each region once in the first year & follow up on issues in the second year
  - Format has been half-day per federation (ROC):
    - Present organization, task assignments to partners, progress & work done
    - Point out problems, successes & issues

- **The exercise has been useful for SA1 management and the federations concerned:**
  - Brought out a number of issues that are activity-wide (project..)
    - WBS – difficulty in adapting a single WBS to many widely different situations
    - Good work “hidden” in the regions – all regions – need to encourage collaboration
- **Observations:**
  - Have looked at 8 regions so far: all are very different!!
    - Countries, sites, languages, organizational structure, …

- **Will be input to the deliverable on ROC status**

Issues for SA1

- **gLite portability**
  - Essential for further uptake, coexistence, interoperability,…
- **Compute Elements – not clear on best way to build production quality CE for the future**
  - EGEE has unprecedented experience at this scale
- **Service manageability issues – must be addressed**
  - Monitoring groups will make some “bolt-on” fixes – but the underlying services must (re-)design this in
- **User support still must improve apparent response times**
- **Usability of the services**
  - UIG, documentation, etc
  - Consistent tools/interfaces
- **Allocation of resources to new VOs**

Plans for 2nd year

- **No changes to plan of work for SA1**

- **Focus on:**
  - Monitoring – for all stakeholders – users, VO managers, site managers, grid operations, project managers. Includes automated publication of standard metrics
  - Upgrade of middleware to support SL4, 64-bit, and other Linuxes – (Debian, SuSe?)
  - Replacement of RBs with gLite WMS in production
  - Replacement of LCG-CE with gLite CE or CREAM
  - User support improvements – all aspects: GGUS process and response, UIG-like activities
    - Important to keep sufficient effort in the ROCs to implement the process
  - SLAs – implement simple SLAs based on reliable metric or monitoring information
  - Interoperability and coexistence remain important
  - Understand long term sustainability issues and organization

Summary

- **Infrastructure has continued to increase in size and scale**
- **Demonstrated sustained workloads – close to 100k jobs/day**
- **Operations processes are significant achievement of the activity**
  - Have evolved during EGEE and EGEE-II
  - Used in related infrastructure projects
  - Basis of interoperations
- **Incremental release process well accepted by sites – no more "big-bang" releases**
- **Interoperability and interoperation is a fact – used in production**
- **User support has improved – but still more to do to improve acceptance**

- **Usability, reliability, manageability … still to be improved**