



Enabling Grids for E-science

SA1 All-Activity Meeting

Ian Bird



Bologna

January 16-17 2007


www.eu-egee.org



PRODUCTION Total number of jobs by VO and DATE
January 2006 - December 2006



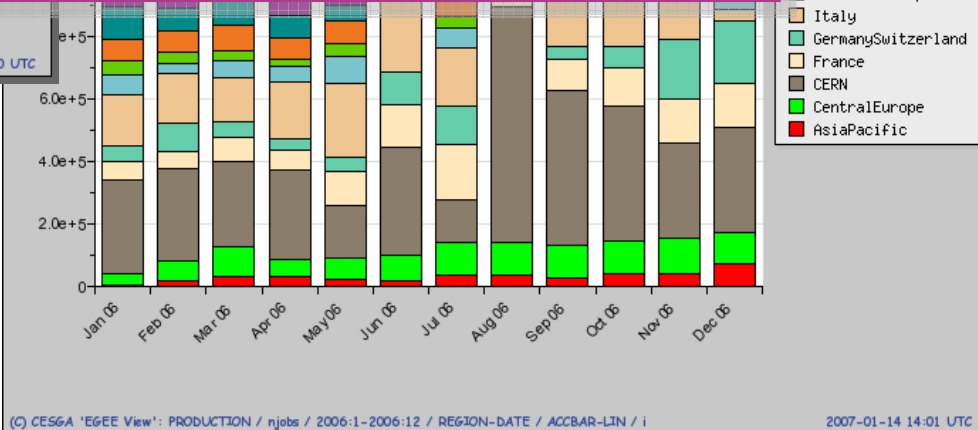

PRODUCTION Total number of jobs by REGION and DATE
January 2006 - December 2006



~17.5 million jobs run (6450 cpu-years) in 2006;
Workloads of the “other VOs” start to be significant – approaching 8-10K jobs per day; and 1000 cpu-months/month

- one year ago this was the overall scale of work for *all* VOs


(C) CESGA 'EGEE View': PRODUCTION / njobs / 2006:1-2006:12 / VO-DATE / ACCBAR-LIN / i 2007-01-14 14:00 UTC




(C) CESGA 'EGEE View': PRODUCTION / njobs / 2006:1-2006:12 / REGION-DATE / ACCBAR-LIN / i 2007-01-14 14:01 UTC

<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	7800	1800	1000*
<i>UK/I</i>	2	23	7309	2010	2410
<i>Italy</i>	1	27	3336	2280	1350
<i>France</i>	1	10	3715	1252	1300*
<i>De/CH</i>	2	13	3237	1852	280*
<i>Northern Europe</i>	6	16	2655	1860	200
<i>SW Europe</i>	2	13	1125	898	200*
<i>SE Europe</i>	8	26	1286	1189	33
<i>Central Europe</i>	7	21	1759	1163	90
<i>Russia</i>	1	15	515	445	52
<i>Asia-Pacific</i>	8	19	1239	751	120
<i>North America</i>	2	8	2973	-	300
Totals	40	192	36949	20265	7335
<i>PPS</i>	14	24	2517	-	-


* Estimates taken from reporting as IS publishes total MSS space





>20k jobs running simultaneously



Use for massive data transfer




Large LHC experiments now transferring ~ 1PB/month each



Tier-1/0 Site Availability : Daily Report

Overall Service Availability for site CNAF : Daily Report






Individual Service Availability for site CNAF : Daily Report

CE

SE

SRM






Service Instance Availability for site CNAF : Daily Report

Site Services

CE : ce01-lcg.cr.cnaf.infn.it

CE : ce03-lcg.cr.cnaf.infn.it

SE : castorgrid.cr.cnaf.infn.it



- **Improved release rollout**
 - Had been lots of complaints (communication, frequency of updates, config changes, planning)
 - Now: (efforts from both SA1 and SA3)
 - Regular updates of repository
 - Announce major changes at ops meeting
 - Announce critical updates with timeline for required updates by sites, otherwise sites do updates when they can
 - Documentation on releases has improved significantly
 - Announcements now have templates and people know where to look for information
 - Remaining issue: to clarify when reconfiguration of service is needed (developers now have to provide this info in Savannah before patch gets accepted)
- **PPS is mature and regular part of process**
 - But see issues ...
- **Grid Operations is now becoming mature**
 - Procedures have evolved, interoperation with OSG main area of development

- **Major problem: (see reviewer comment on porting!)**
 - With the current state of gLite, SA1 is not able to provide and maintain compatibility with the physical deployment environment.
 - There is no port to SLC4 – let alone the promised other OS's
 - This means we cannot even maintain operations at our current sites as they need to upgrade their hardware
 - Seems to be due to the complexity of the gLite packages, interdependencies, external dependencies and consequent problems to build and port
 - This has been raised as an issue for years but nothing has been done about it
 - This is now a crisis in the making – it is both a management and a technical problem

 - SA1 must do something about this if the middleware providers will not
 - We have committed to providing a production infrastructure – we are not able to maintain it given the current state
 - gLite middleware is too heavyweight and inflexible, and very hard to deploy
 - What are we going to do about it?

- **Pre-production service: Is it value for money?**
 - The PPS has been used:
 - During gLite-3.0 deployment/testing
 - For major changes (eg SLC4 testing)
 - By certain VOs (SWITCH, DILIGENT) as a main platform
 - But by and large it is not heavily used
 - Use is rather intermittent, and mainly for small tests – it cannot reproduce the environment and scale of the LCG VOs production environments
 - EIS team (and ATLAS) use to check that bugs are fixed
 - To make it available for when it is needed it has to be kept up to date
 - Weekly release cycle
 - Well maintained and good response from the sites involved
 - All of the sites involved have volunteered and want to participate (~20 sites)
 - Perhaps the scale can be reduced to be more appropriate for the actual usage?
 - Use more as extended deployment testbed / extended certification (with SA3)?
 - Cover major deployment scenarios
 - Or, can it be made more usable by, more useful to, the VOs?
 - Need NA4 to help us answer this

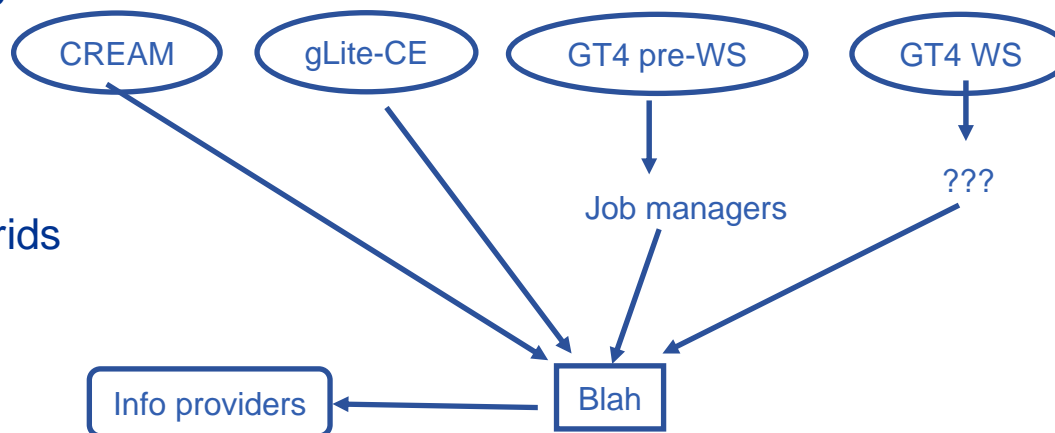
- **gLite CE breaks interoperability with other infrastructures, sites , and applications**
 - Only common interface will be through Condor-G
- **Lack of management ability of grid services**
 - Remains a major issue for site and service managers
 - Very little has changed yet
- **Web site and information, &**
- **UIG –**
 - Still see the need for help in getting users started
 - Not yet a framework even?
- **Resources for new VOs!**
 - See EGEE-3 discussion; this needs to be addressed by the resource owners

- **Preparation for LHC start-up**
 - This will be the main driving force behind how the infrastructure changes
 - Dress rehearsals start in June; expect accelerator start-up in November
 - Crucial that the infrastructure is ready, reliable, and stable
 - Must ensure that changes are not disruptive – probably even more important than new functionality
- **Middleware changes anticipated in deployment**
 - Migration to SLC4 + 64 bit support
 - SRMv2.2 support: In the Storage Elements, but also in the DM tools (GFAL, lcg-utils, FTS)
 - Move to the gLite WMS: but this is still not yet good enough to be deployed widely as a production service
 - Implementation and deployment of the gLite-CE
 - But this has its own problems – see slide
 - Replacement of gLiteIO/Fireman with LCF/GFAL (April?)
 - glxec provided as a tool – sites will decide how it can be used
 - Job priorities – following wg proposals

Improvement of service reliability – also to respond to reviewers comments on understanding of the overall system and job failures

- **SAM – continue to develop – need to complete the set of tests**
 - Will publish availability/reliability of all sites (currently done for LCG Tier 1s)
- **Three working groups have been set up to focus on improving service and site reliability – all 3 are coordinated together**
 - Grid Monitoring
 - Pull together existing monitoring data, provide views for stakeholders
 - Not more development, but bringing together what exists, and fill in the gaps
 - Early deliverable – integrate remote monitors of a service into site fabric monitoring
 - Site Management
 - Set up through HEPiX, but relevant to all system admins.
 - Tools, best practices, recommendations, for improving site management
 - *E.g. recommend a set of tools to use for new sites*
 - System Analysis
 - Continue the work done by ARDA on experiment dashboards – but coordinated with the other groups
 - Provides a level of understanding from app p.o.v and can feedback to ops and other monitors
- **Need coordinated debugging with applications to find underlying failures – and fix them**

- Issue: gLite CE/Cream does not support all the submission interfaces we have now
 - If we move to deploy the gLite CE only – will force other VOs to use Condor as *the standard* interface to CEs
 - Alternatives:
 - Accept and push Condor as the interface
 - Deploy an architecture similar to that proposed here
 - Risk if we do not do this:
 - Break support for existing apps
 - Break co-existence with other grids at sites
 - What if national/campus grids deploy GT4???
- Blah must be the **ONLY** interface to the batch systems
 - Logging, accounting, info providers, ...
 - Cream must be added as a plugin to this architecture (i.e. no shortcuts around blah)



- **Complete 1st round of federation/partner reviews**
 - 1st set (CE, SEE) were done in November
 - 2nd set (DE/CH, NE, SWE, UKI) in February
- **GGUS:**
 - Workshops with VOs to improve user friendliness
 - 1st round with LCG VOs next week
 - 2nd round with other VOs planned for ???
- **Operations workshops**
 - Next week during wLCG collaboration meeting
 - June in US focus on interoperation
 - ?May? in Europe (default CERN) general operations workshop

- **17. Investigate how the VO resource access policies relate to the above issues and objectives.**
 - This is outstanding – we have no good response
- **18. Data privacy issues may arise if the accounting data are not managed carefully. Appropriate access control policies should be explored to achieve that objective.**
 - Handled by accounting system

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