



Enabling Grids for E-science


SA1: Grid Operations and Management

Ian Bird, CERN
SA1 Activity Manager
EGEE 2nd EU Review
6-7/12/2005

www.eu-egee.org



- **Scale and usage of infrastructure**
- **Grid Operations**
 - Metrics, operations support
- **Pre-production Service**
- **User support**
- **Operational security**
- **Interoperability / interoperation**
 - Input to standards process
- **Certification and deployment process**
- **gLite certification**
 - LCG-2/gLite convergence
- **Key points for SA1**
- **Plans for next period**






EGEE Grid Sites : November 2005



- Many more sites than anticipated for this stage of project
 - 179 actual, cf. 50 proposed for end of year 2
- Includes industrial partner sites (HP in Puerto Rico and UK)
- Exposes full complexity of grid operations - no. sites not CPU

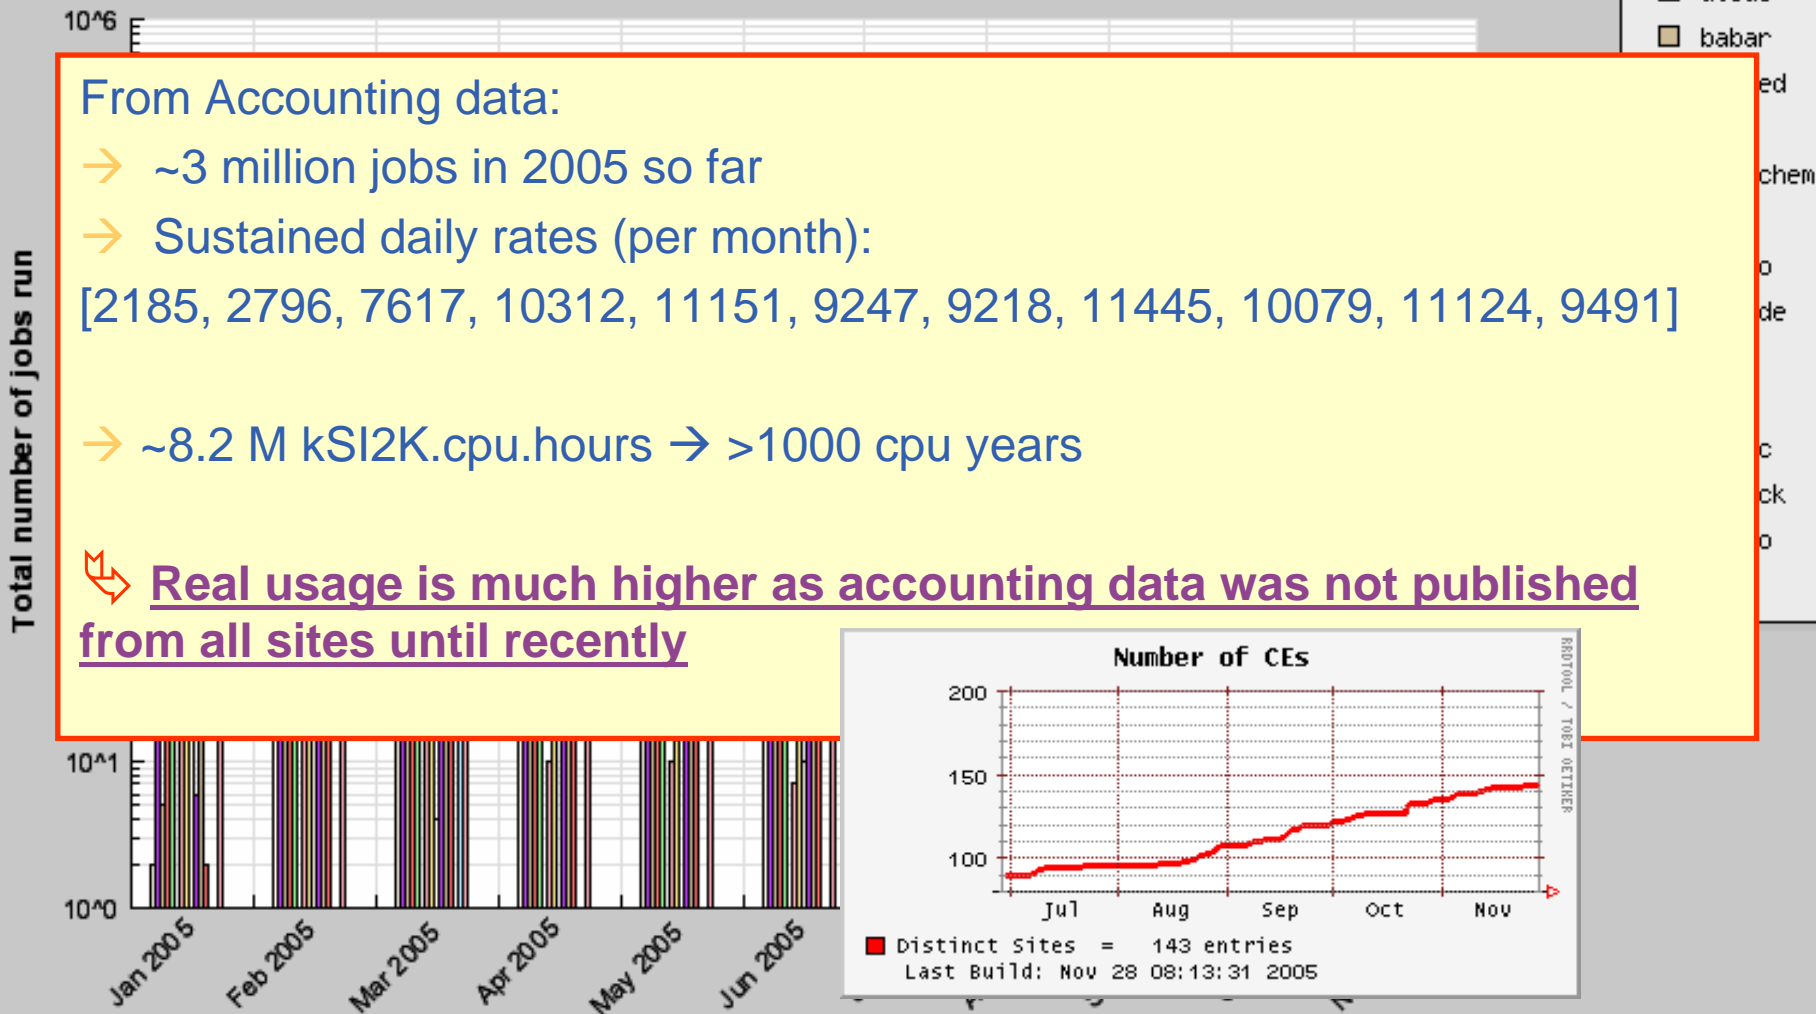
EGEE:



179 sites, 39 countries
 >17,000 processors,
 ~5 PB storage



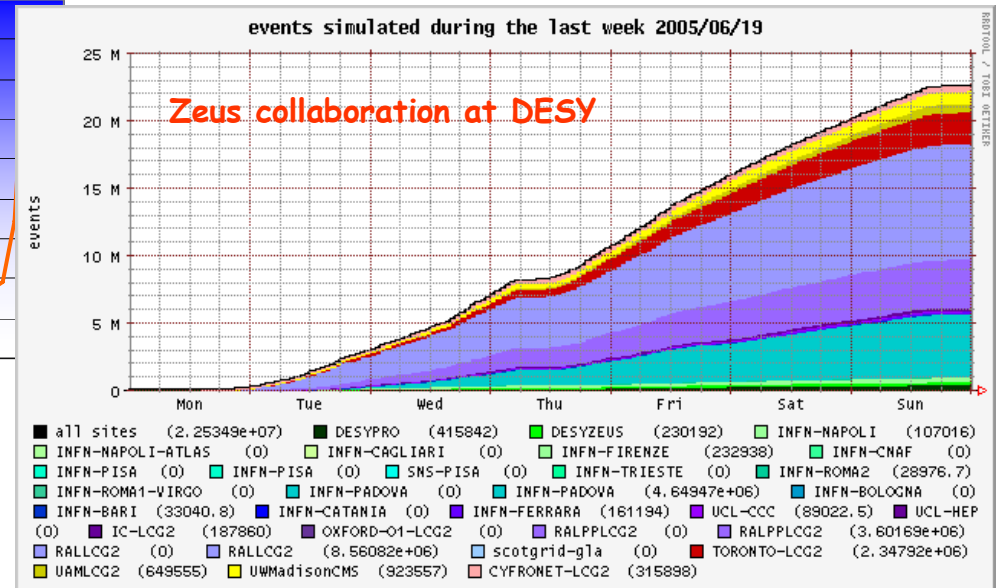
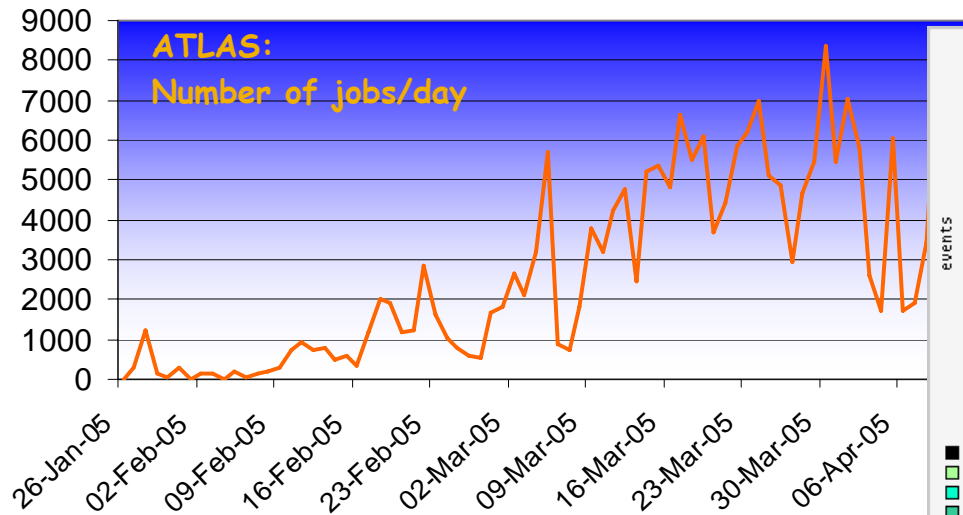
country	sites	country	sites	country	sites
Austria	2	India	2	Russia	12
Belgium	3	Israel	3	Singapore	1
Bulgaria	4	Italy	25	Slovakia	4
Canada	7	Japan	1	Slovenia	1
China	3	Korea	1	Spain	13
Croatia	1	Netherlands	3	Sweden	4
Cyprus	1	Macedonia	1	Switzerland	2
Czech Republic	2	Pakistan	2	Taiwan	4
Denmark	1	Poland	5	Turkey	1
France	8	Portugal	1	UK & Ireland	37
Germany	10	Puerto Rico	1	USA	4
Greece	6	Romania	1	Yugoslavia	1
Hungary	1				


Aggregate Accounting Plot for EGEE



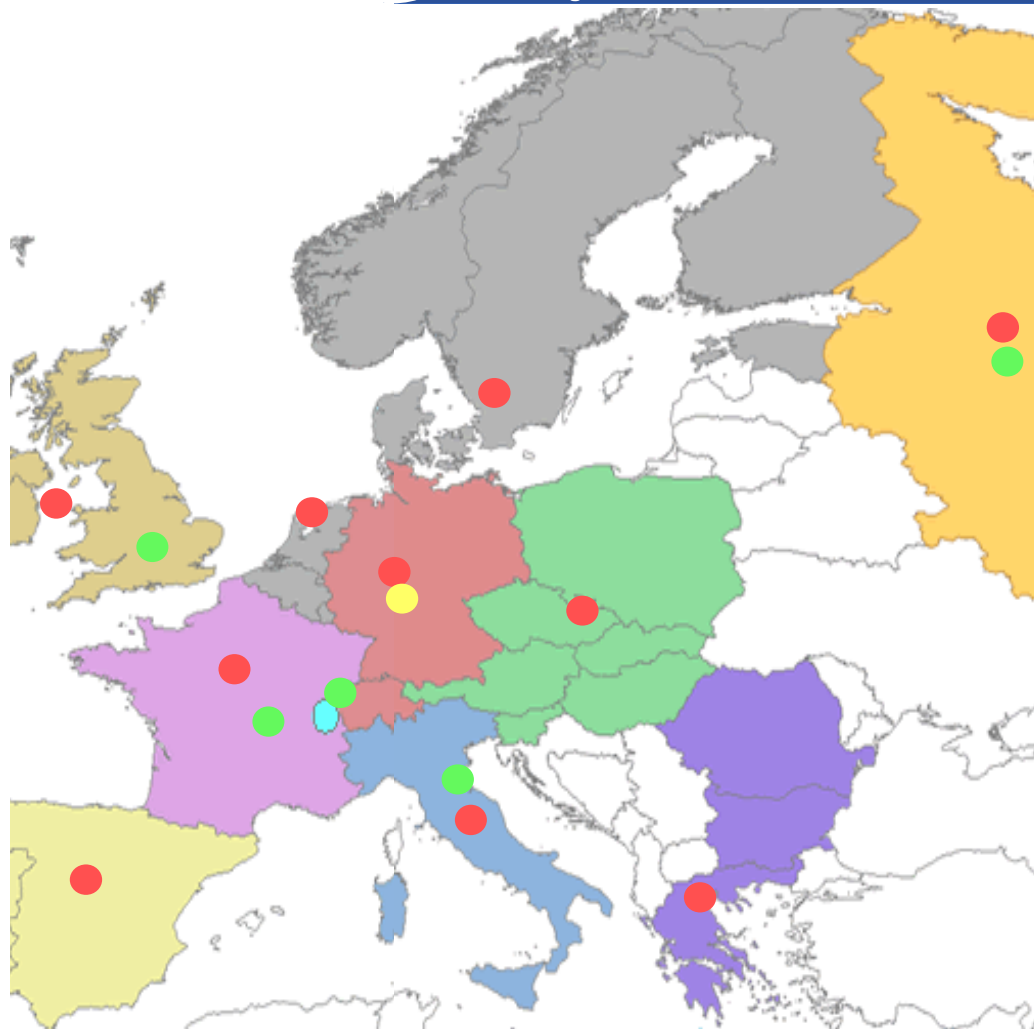


LCG sustained data transfers using FTS; in excess of 500 MB/s



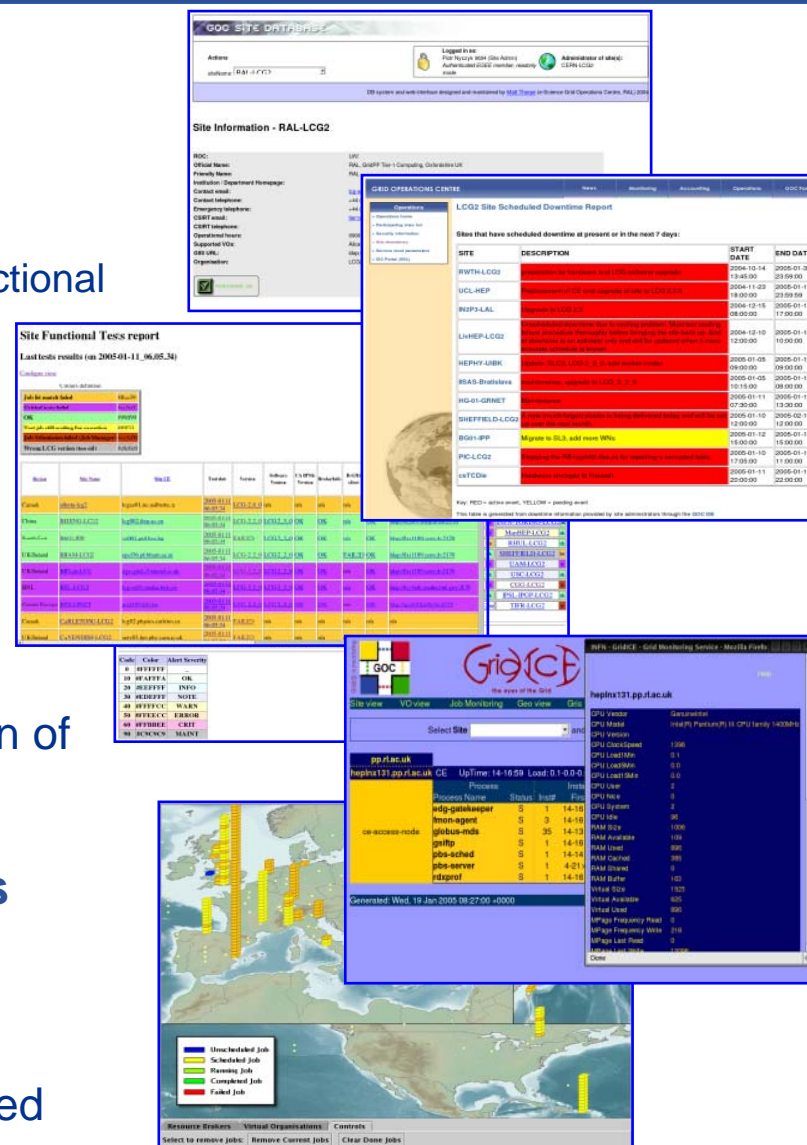


- SC2** – Reliable data transfer (disk-network-disk) – 5 Tier-1s, aggregate 500 MB/sec sustained at CERN
- SC3** – Reliable base service – most Tier-1s, some Tier-2s – basic experiment software chain – grid data throughput 500 MB/sec, including mass storage (~25% of the nominal final throughput for the proton period)
- SC4** – All Tier-1s, major Tier-2s – capable of supporting full experiment software chain inc. analysis – sustain nominal final grid data throughput
- LHC Service in Operation** – September 2006 – ramp up to full operational capacity by April 2007 – capable of handling twice the nominal data throughput



- **Operations Management Centre (OMC):**
 - At CERN – coordination etc
- **Core Infrastructure Centres (CIC)**
 - Manage daily grid operations – oversight, troubleshooting
 - “Operator on Duty”
 - Run essential infrastructure services
 - Provide 2nd level support to ROCs
 - UK/I, Fr, It, CERN, Russia, Taipei
- **Regional Operations Centres (ROC)**
 - 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

- **Operator on duty**
 - Started November 2004
 - Crucial in stabilising sites
 - Many complementary monitoring tools
 - Essential tools : GIS monitor and Site Functional Tests
- **Simplified VO selection of good sites**
 - VO can select set of functional tests that it requires
 - Can white- or black-list sites
 - Can include VO-specific tests (e.g. sw environment)
 - SFT framework provides dynamic selection of “good” sites
- **SFT’s have evolved to become stricter as lessons are learned**
- **Normally >80% of sites pass SFTs**
 - NB of 180 sites, some are not well managed

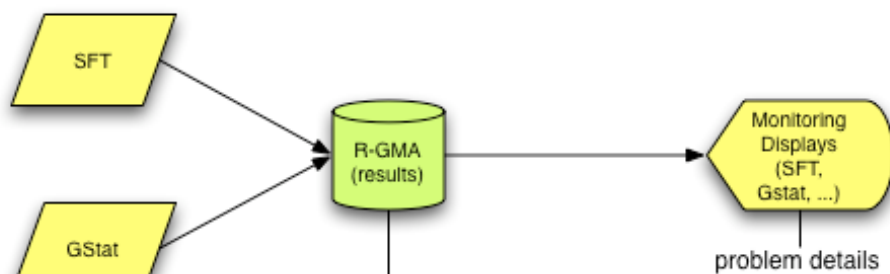


The collage includes several key interfaces:

- GOC SITE DATABASE:** A web-based portal for site information, showing details for RAL-LCG2.
- GRID OPERATIONS CENTRE:** A dashboard with navigation tabs for Operations, Monitoring, Accounting, Operations, and GOC Portal.
- LCG2 Site Scheduled Downtime Report:** A table listing scheduled downtime for various sites.


SITE	DESCRIPTION	START DATE	END DATE
WUFI-LCG2	Hardware maintenance (CPU, RAM, network)	2004-10-14	2005-01-31
UCL-HEP	Hardware maintenance (CPU, RAM, network)	2004-11-23	2005-01-17
RDUP-LAL	Hardware maintenance (CPU, RAM, network)	2004-12-16	2005-01-17
LHNER-LCG2	Hardware maintenance (CPU, RAM, network)	2004-12-10	2005-01-18
HEPHY-LARK	Hardware maintenance (CPU, RAM, network)	2005-01-08	2005-01-13
IBAS-Brunel	Hardware maintenance (CPU, RAM, network)	2005-01-05	2005-01-12
HQ-41-GRNET	Hardware maintenance (CPU, RAM, network)	2005-01-11	2005-01-11
SHEFFIELD-LCG2	Hardware maintenance (CPU, RAM, network)	2005-01-10	2005-02-10
BGI-IPP	Upgrade to G.L.S. web server	2005-01-12	2005-01-14
PK-LCG2	Upgrade to G.L.S. web server	18.00.00	18.00.00
ISTEDN	Hardware maintenance (CPU, RAM, network)	2005-01-11	2005-01-11
ISTEDN	Hardware maintenance (CPU, RAM, network)	2005-01-11	2005-01-11
- Site Functional Test report:** A detailed report for the last tests results (on 2005-01-11_06:05:34), showing a grid of test results for various sites.
- GridCE Monitoring Service:** A web interface for monitoring the GridCE service, showing site selection and monitoring options.
- Job Monitoring:** A detailed view of a job (hepin131.pp.rl.ac.uk) showing its status, processes, and resource usage.
- Map:** A geographical map showing the locations of various sites across Europe and Asia, color-coded by job status.

- **Weekly operations meetings**
- **Regular ROC, CIC managers meetings**
- **Series of EGEE Operations Workshops**
 - Nov 04, May 05, Sep 05
 - Last one was a joint workshop with Open Science Grid
- **These have been extremely useful**
 - Will continue in Phase II
 - Bring in related infrastructure projects – coordination point
 - Continue to arrange joint workshops with OSG (and others?)




- **CIC-on-duty: currently 6 teams (CERN, IN2P3, RAL, INFN, Russia, Taipei) working in weekly shifts**
- **The operators look at emerging**

- **Procedures described in detail in the Operations Manual**
- **Geographically distributed responsibility for operation: there is no “centre”; Tools are hosted at different sites:**
 - **GOC DB (UK), GStat (Taipei), SFT (CERN), CIC portal (Lyon)**

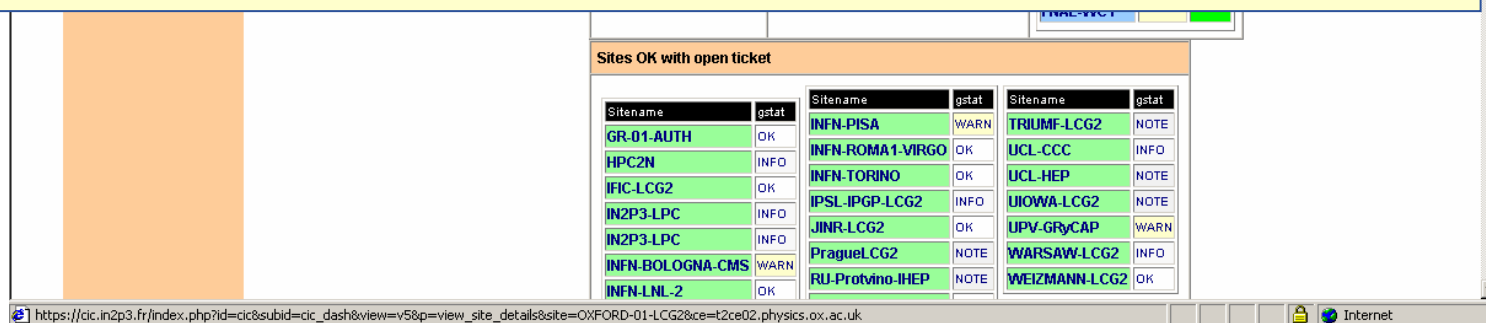


- **ROC is responsible for timely problem solution - otherwise ticket is escalated**
- **Priorities and deadlines for tickets are set depending on site size (number of CPUs)**



Problem categories

- **Main tool for CIC-on-duty**
 - Integrated view of monitoring tools (summary) - shows only failures and assigned tickets
 - Detailed site view with table of open tickets and links to monitoring results
 - Single tool for ticket creation and notification emails with detailed problem categorisation and templates
 - Ticket browser with highlighting expired tickets
- **Well maintained - adapts quickly to new requirements/suggestions**



Sitename	gstat	Sitename	gstat	Sitename	gstat
GR-01-AUTH	OK	INFN-PISA	WARN	TRIUMF-LCG2	NOTE
HPC2N	INFO	INFN-ROMA1-VIRGO	OK	UCL-CCC	INFO
IFIC-LCG2	OK	INFN-TORINO	OK	UCL-HEP	NOTE
IN2P3-LPC	INFO	IPSL-IPGP-LCG2	INFO	UIOWA-LCG2	NOTE
IN2P3-LPC	INFO	JINR-LCG2	OK	UPV-GRyCAP	WARN
INFN-BOLOGNA-CMS	WARN	PragueLCG2	NOTE	WARSAW-LCG2	INFO
INFN-LNL-2	OK	RU-Protvino-IHEP	NOTE	WEIZMANN-LCG2	OK

- Shows results matrix with all sites
- Detailed test log available for troubleshooting debugging
- History of information is kept
- Selection of “critical” tests for each VO to define which sites are good/bad
- “Freedom of Choice” tool:
 - Allows apps to select good sites according to their criteria

Very important in stabilising sites:

- Apps use only good sites
- Bad sites are automatically excluded
- Sites work hard to fix problems

Test summary		Colours definition		Test abbreviations																
SD	Scheduled downtime	#a3a3a3		cs	CSH test															
JL	Job list match failed	#aab3ff		swdir	VO software directory															
JS	Job submission failed	#f4876b		rgma	R-GMA															
CT	Critical tests failed	#f9d48e		dirac-test	Dirac full test															
				ver	Software Version (WN)															
				wn	WN host name															
				ca	CA certs version															
				crl	CRL timestamp test															
				rm	Replica Management															
				votag	VO Tag management															
				js	Job submission															
				brkr	BrokerInfo															
VO lhcb																				
				rm	votag	swdir	crl	St.	js	dirac-test										
				X	O	O	!!!	JL	X	??										
				O	O	O	!!!	JL	X	??										
3.	OK	TOKYO-LCG2	dgce0.icepp.jp	OK	O	2	4	0	I	O	O	O	O	O	O	!!!	JL	X	??	
4.	OK	Taiwan-LCG2	lcg00125.grid.sinica.edu.tw	OK	O	2	6	0	I	O	O	O	O	O	O	!!!	JL	X	??	
5.	OK	Taiwan-IPAS-LCG2	testbed001.phys.sinica.edu.tw	OK	O	2	6	0	I	O	O	O	O	O	O	!!!	JL	X	??	
6.	OK	GOG-Singapore	mclon.ngpp.ngp.org.sg	OK	O	2	6	0	I	O	O	O	O	O	O	!!!	JL	X	??	
7.	OK	Taiwan-NCUCC-LCG2	ce.cc.ncu.edu.tw	OK	O	2	6	0	I	O	O	O	O	O	O	!!!	OK	O	O	
8.	OK	LCG_KNU	cluster50.knu.ac.kr	OK	O	2	5	0	I	O	O	O	O	O	O	!!!	CT	O	!!!	
BNL																				
9.	SD	BNL-LCG2	lcg-ce01.usatlas.bnl.gov	SD	X	??	??	??	??	??	??	??	??	??	??	??	SD	X	??	
Canada																				
10.	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	

- Monitoring tool for Information System:
 - Periodically queries all Site BDII's (but not Top-level BDII's)
 - Checks if Site BDII's are available
 - Checks integrity of published information
 - Checks for missing entities, attributes
 - Detects and reports information about some of the Services: RB, MyProxy, LFC but doesn't monitor them
 - Detects duplicated services in some cases (eg. 2 global LFC servers a single VO)

GStat: 12:48:18 11/08/05 GMT

[home](#) [alert](#) [table](#) [service](#) [regional](#) [service](#) [metrics](#) [links](#) [?](#) [prod](#) [test](#) [seegrid](#)

CERN-PROD Status: OK OK GOCD graphs

GOCD Configuration information:
 status: Certified, type: Production
 giis url: ldap://ce02-slc3.cern.ch:2170/mds-vo-name=CERN-PROD,o=grid

To test site GIIS: ldapsearch -x -H ldap://ce02-slc3.cern.ch:2170 -b mds-vo-name=CERN-PROD,c


BDII Node Check: alert history [?](#)

CERN SE Check: alert history [?](#)

No BDII Node to check in GOCD
 test: ldapsearch -xLLL -l 15 -h -p 2170 -b 'GlueSEUniqueID=lxn1183.cern.ch,mds-vo-name=cer

GIIS Perf Check: ok alert history [?](#)

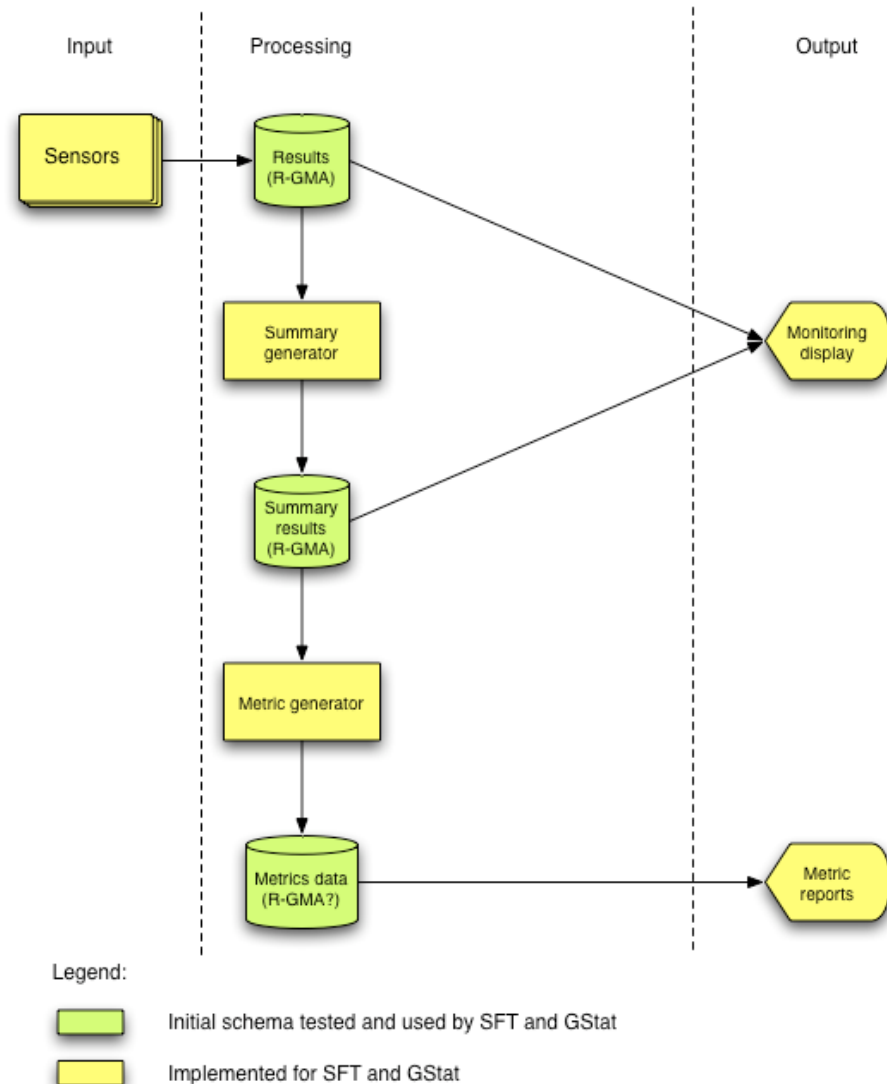
Query Response Time (ms): 2594.3 - OK
 GIIS Entries Found: 29 - OK
 GIIS Old Entries Found: 0 - OK




GIIS Sanity Check: ok alert history [?](#)

Passed

- Using current data schema and R-GMA - integrate monitoring information from SFT and GStat
- Summary generator uses list of critical tests to generate a summary per site - binary value (good/bad) generated every 1h
- Metric generator integrates the summaries over time period (1 day...) to generate availability metric



Evolution of SFT metric



<i>Service</i>	<i>Class</i>	<i>Comment</i>
SRM 2.1	C	Monitoring of SE
LFC	C/H	
FTS	C	Base on SC experience
CE	C	Monitored by SFT now
RB	C	Job monitor exists
Top level BDII	C	Can be included in Gstat
Site BDII	H	Monitored by Gstat
MyProxy	C	
VOMS	H	
R-GMA	H	

Effort identified for each service

Will all be integrated into SFT framework

First approach to SLA:

- each Class (C, H, etc) defines required service availability

Checklist for a new service

- **User support procedures (GGUS)**
 - Troubleshooting guides + FAQs
 - User guides
- **Operations Team Training**
 - Site admins
 - CIC personnel
 - GGUS personnel
- **Monitoring**
 - [unclear]
 - [unclear]
- **Accounting**
 - [unclear]
- **Service Parameters**
 - Scope - Global/Local/Regional
 - SLAs
 - Impact of service outage
 - Security implications
- **Contact Info**
 - Developers
 - Support Contact
 - Escalation procedure to developers
- **Interoperation**
 - Documented issues
- **First level support procedures**
 - How to start/stop/restart service
 - How to check it's up
 - Which logs are useful to send to CIC/Developers
 - and where they are
- **Tools for CIC to spot problems**
 - GIS monitor validation rules (e.g. only one "global" component)
 - Definition of normal behaviour
 - Metrics
- **CIC Dashboard**
 - Alarms
- **Deployment Info**
 - RPM list
 - Configuration details (for yaim)
 - Security audit

➤ **What is now understood as essential to make a reliable production service from a middleware component**

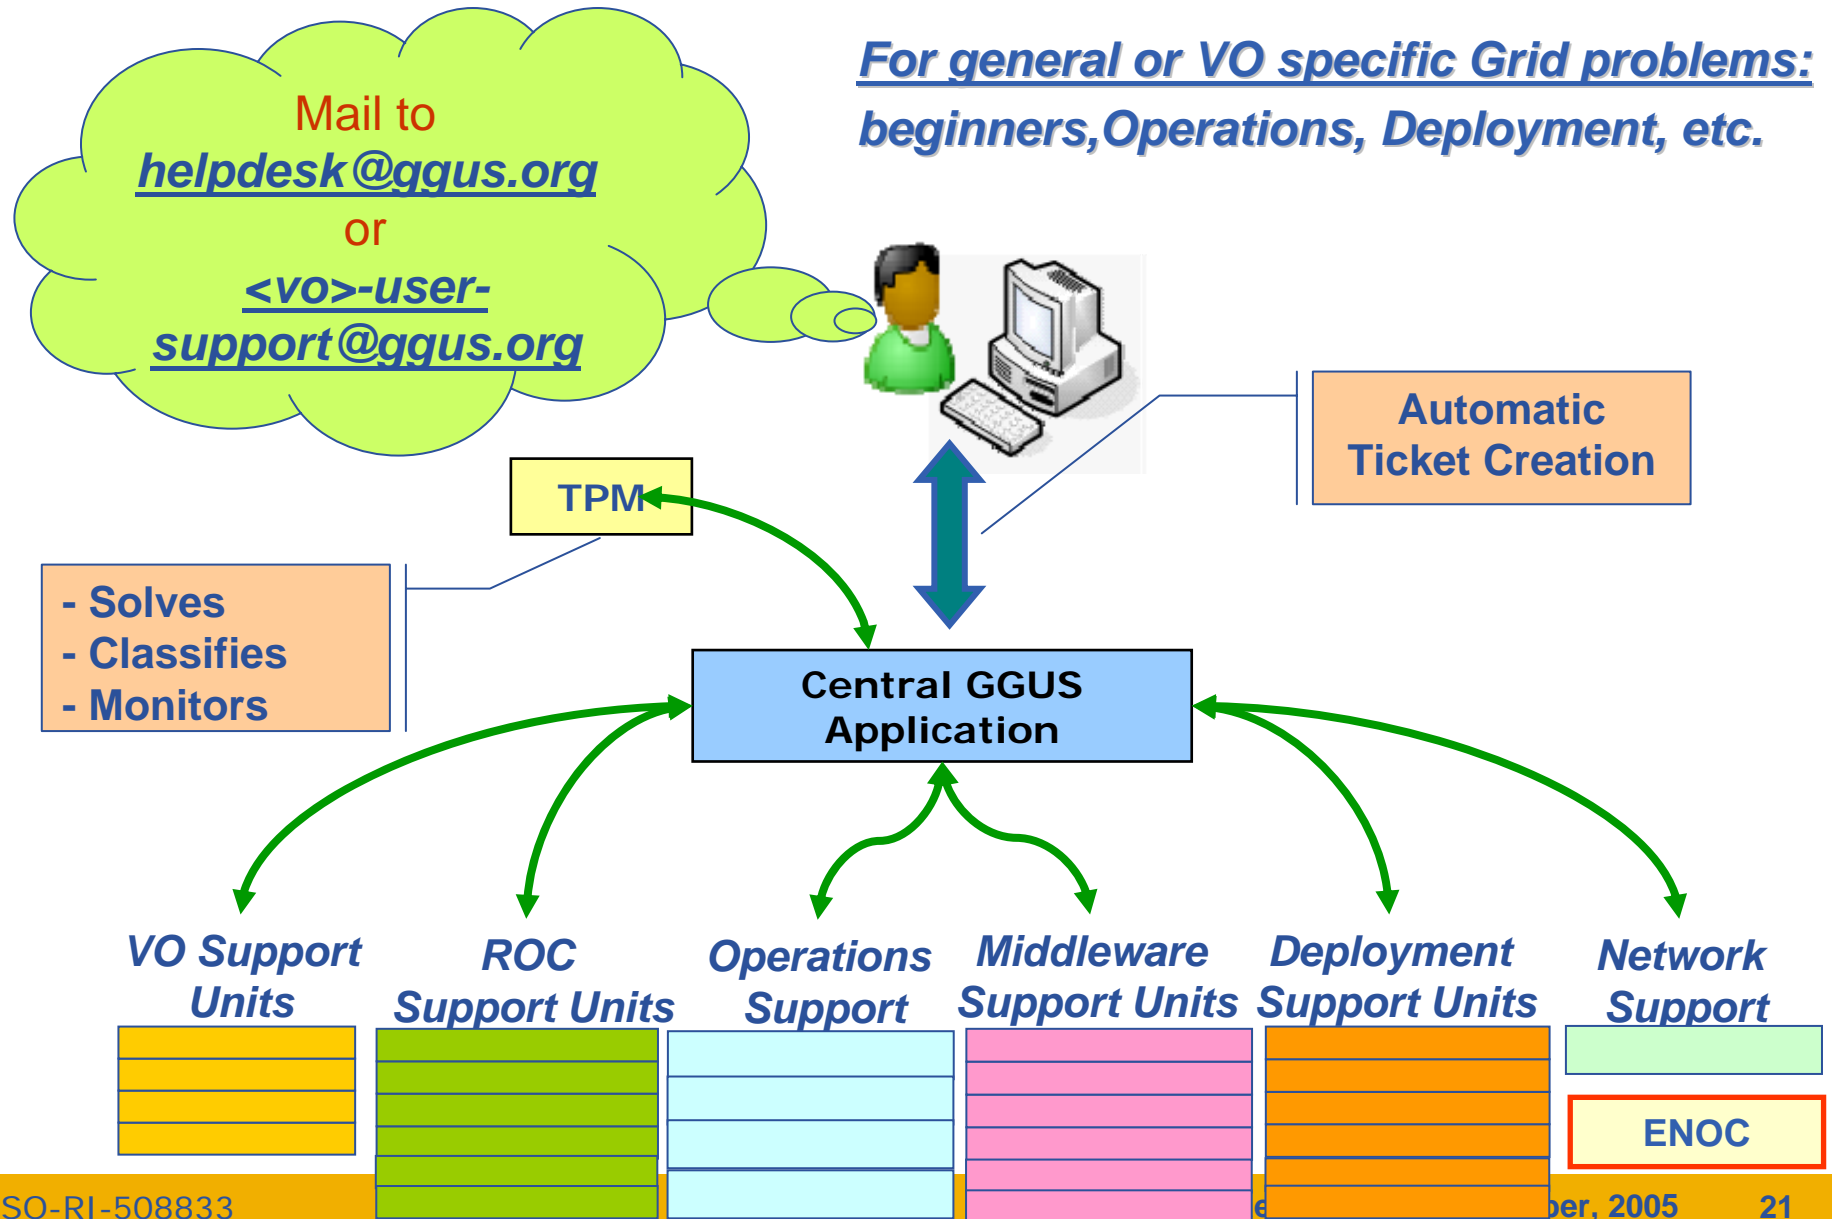
➤ **Not much middleware comes with all this ...**

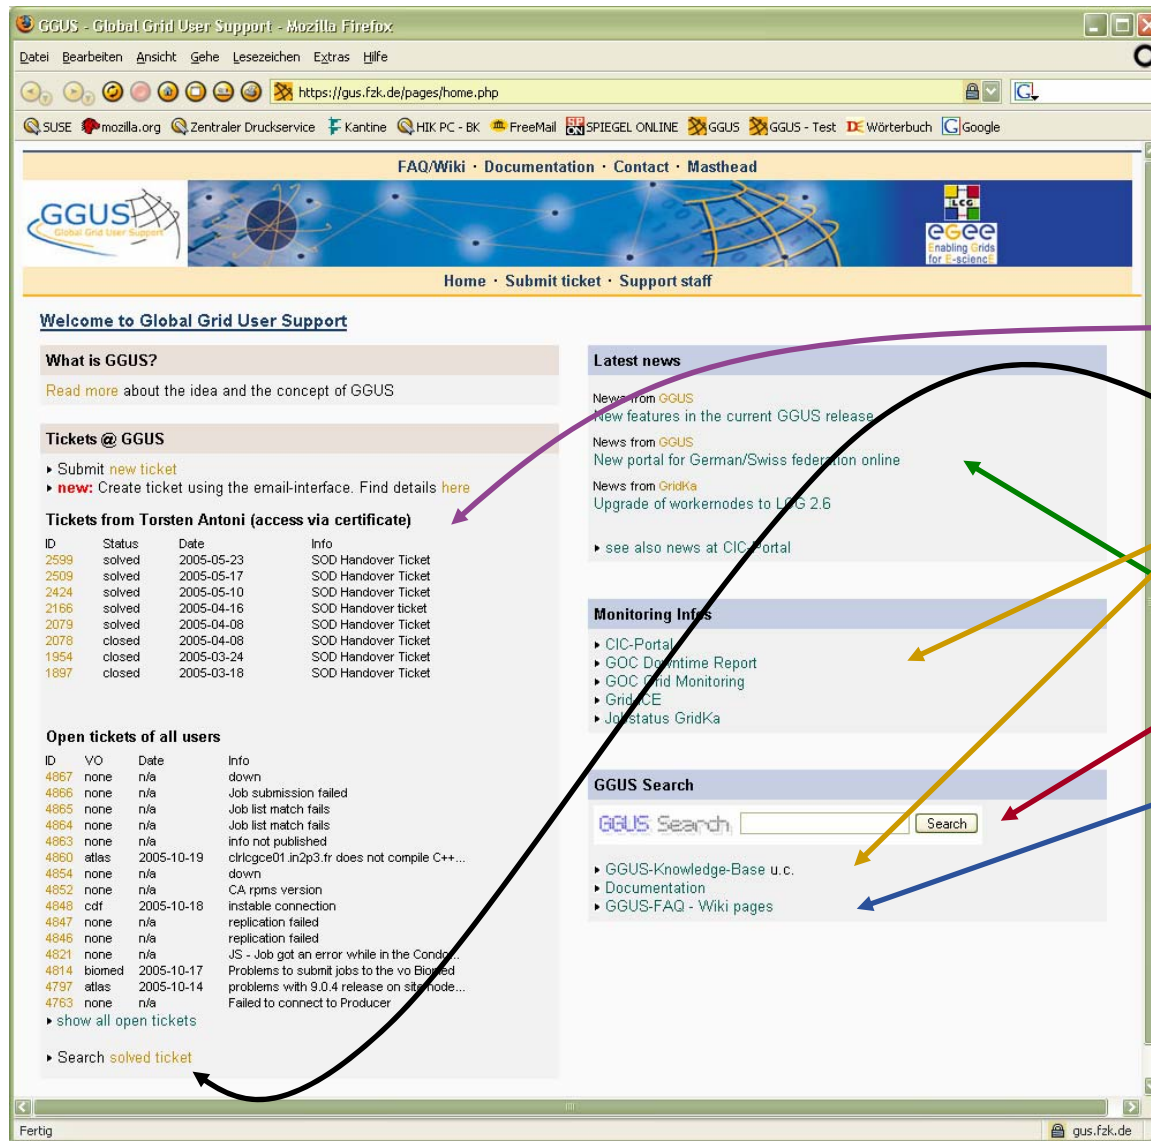
- **Current PPS is a “pure” gLite service**
 - BDII, SRM SE and MyProxy server are also needed.
- **The PPS is available and used by many VOs**
 - HEP VOs (CMS, ATLAS, Alice, LHCb)
 - ARDA
 - BioMed
 - egeode
 - NA4 (testing)
 - DILIGENT
 - SWITCH
- **Currently upgrading from gLite 1.4 to gLite 1.4.1 (a major patch)**
 - As the service is now in use, upgrades are planned and phased to minimize the impact to users.
- **Currently preparing to move the day-to-day operations of the PPS to the production operations team**
 - SFT monitoring is now in place
 - All PPS sites are now correctly entered in the GOC DATABASE
 - Production operations processes are being implemented for the PPS (so far the CA upgrade process and scheduled downtime processes are in place).
- **Planning is under way for moving the PPS from being a pure gLite service to being a true pre-production service which closely mirrors production (+ new and updated functionality and services)**

ROC	Site	CPUs	SE	Core Services			
Asia-Pacific	ASGC	?		WMS			
CE	CYFRONET	3					
CERN	CERN	54	DPM	WMS	FTS	VOMS (production)	
DE/CH	FZK	2					
France	IN2P3	?			FTS	VOMS	
Italy	CNAF	150	DPM	WMS		VOMS	BDII
Italy	INFN-Padova	?					
NE	NIKHEF	0				VOMS	
SEE	UoM	2					
SEE	UPATRAS	3		WMS			
SWE	CESGA	2					R-GMA
SWE	IFIC	1	Castor				
SWE	LIP	2	DPM				MyProxy
SWE	PIC	180	Castor	WMS			FireMan
UK/I	ScotGrid-Glasgow	0			FTS		

- **PIC, CNAF and CERN have given access to production batch farms**
 - PIC and CNAF running LCG WNs; CERN running gLite WNs.
 - Queue to production batch farm is currently restricted to 50 jobs. This restriction can be removed, increasing the number of CPUs at CERN to ~1,500.
- **To date, over 1.5 million jobs have been submitted to the PPS WMSs.**

- **User Support in EGEE (helpdesk, call-centre)**
 - Regional support with central coordination (GGUS @ FZK)
 - GGUS platform connects:
 - CICs, ROCs, VOs, service teams providing support
 - Middleware developers and support
 - Networking activities (training etc).
 - TPM – oversee problem lifecycle
 - Ensure problems assigned and followed up
 - Problem resolution by volunteer experts – harness informal processes
 - Users can report via local helpdesks, ROC helpdesk, VO helpdesk, or to GGUS
 - Ticket traffic increasing
 - Now: Change in users from a few, experienced, production managers to general users (low quality of tickets)
- **VO support**
 - Other aspect of user support – direct support to apps to integrate with grid middleware
 - Application driven process: set up several task forces to implement this (follow successful model in LCG)





GGUS - Global Grid User Support - Mozilla Firefox

https://gus.fzk.de/pages/home.php

FAQ/Wiki · Documentation · Contact · Masthead

GGUS Global Grid User Support

Home · Submit ticket · Support staff

Welcome to Global Grid User Support

What is GGUS?
Read more about the idea and the concept of GGUS

Tickets @ GGUS
Submit new ticket
new: Create ticket using the email-interface. Find details here

Tickets from Torsten Antoni (access via certificate)

ID	Status	Date	Info
2599	solved	2005-05-23	SOD Handover Ticket
2509	solved	2005-05-17	SOD Handover Ticket
2424	solved	2005-05-10	SOD Handover Ticket
2166	solved	2005-04-16	SOD Handover ticket
2079	solved	2005-04-08	SOD Handover Ticket
2078	closed	2005-04-08	SOD Handover Ticket
1954	closed	2005-03-24	SOD Handover Ticket
1897	closed	2005-03-18	SOD Handover Ticket

Open tickets of all users

ID	VO	Date	Info
4867	none	n/a	down
4866	none	n/a	Job submission failed
4865	none	n/a	Job list match fails
4864	none	n/a	Job list match fails
4863	none	n/a	info not published
4860	atlas	2005-10-19	chrlgce01.in2p3.fr does not compile C++...
4854	none	n/a	down
4852	none	n/a	CA rpms version
4848	cdf	2005-10-18	instable connection
4847	none	n/a	replication failed
4846	none	n/a	replication failed
4821	none	n/a	JS - Job got an error while in the Cond...
4814	biomed	2005-10-17	Problems to submit jobs to the vo Biomed...
4797	atlas	2005-10-14	problems with 9.0.4 release on site node...
4763	none	n/a	Failed to connect to Producer

show all open tickets

Search solved ticket

Latest news
News from GGUS
New features in the current GGUS release
News from GGUS
New portal for German/Swiss federation online
News from GridKa
Upgrade of workmodes to LUG 2.6
see also news at CIC Portal

Monitoring Infos
CIC-Portal
GOC Downtime Report
GOC Grid Monitoring
GridICE
Jobstatus GridKa

GGUS Search
GGUS Search: Search
GGUS-Knowledge-Base u.c.
Documentation
GGUS-FAQ - Wiki pages

Browseable tickets

Search through solved tickets


Useful links (Wiki FAQ)

Latest News

GGUS Search Engine

Updated documentation (Wiki FAQ)

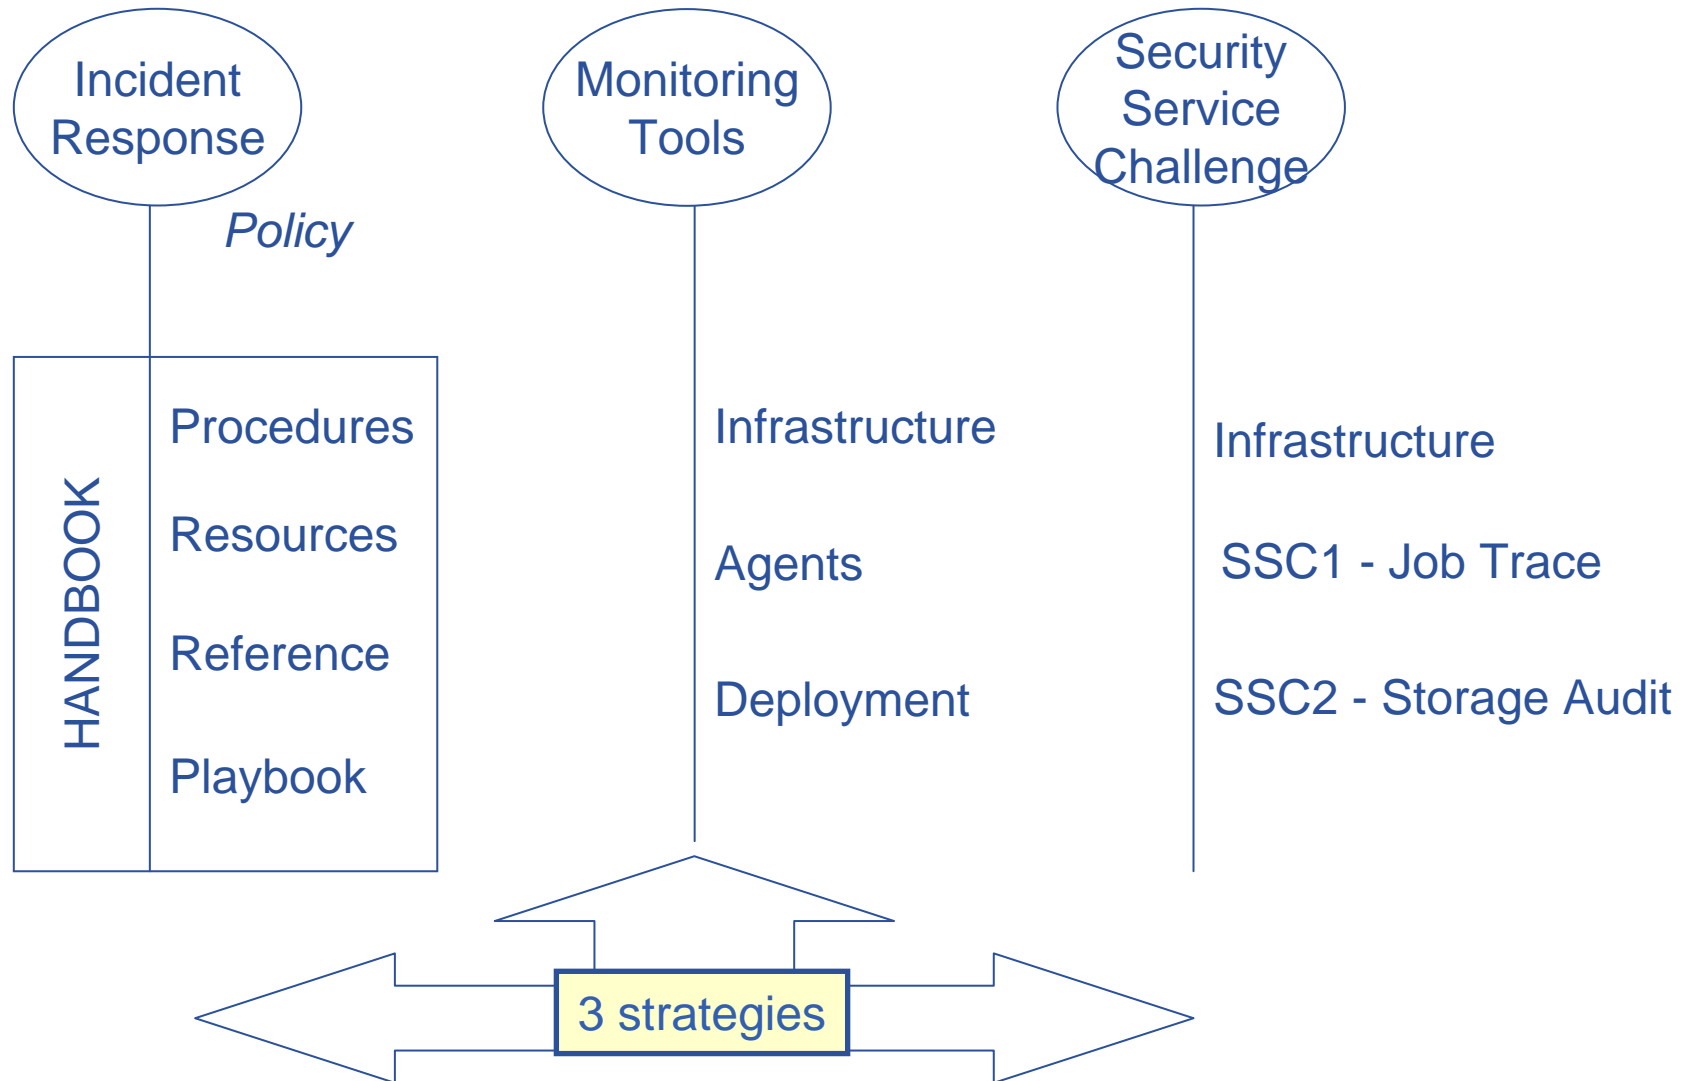
- **Joint Security Policy Group**
 - EGEE with strong input from OSG
 - Policy Set:



- **Policy Revision In Progress/Completed**

- Grid Acceptable Use
 - <https://edms.cern.ch/document/428036/>
 - common, general and simple AUP
 - for all VO members using many Grid infrastructures
 - *EGEE, OSG, SEE-GRID, DEISA, national Grids...*
- VO Security
 - <https://edms.cern.ch/document/573348/>
 - responsibilities for VO managers and members
 - VO AUP to tie members to Grid AUP accepted at registration
- Incident Handling and Response
 - <https://edms.cern.ch/document/428035/>
 - defines basic communications paths
 - defines requirements (MUSTs) for IR
 - *reporting*
 - *response*
 - *protection of data*
 - *analysis*
 - not to replace or interfere with local response plans

- **OSCT membership → EGEE ROC security contacts**
 - What it is not:
 - Not focused on middleware security architecture
 - Not focused on vulnerabilities (see *Vulnerabilities Group*)
 - Focus on Incident Response Coordination
 - Assume it's broken, how do we respond?
 - Planning and Tracking
 - Focus on 'Best Practice'
 - Advice
 - Monitoring
 - Analysis
 - Coordinators for each EGEE ROC
 - plus OSG LCG Tier 1 + Taipei



- **Has been set up this summer (CCLRC lead)**
- **Purpose: inform developers, operations, site managers of vulnerabilities as they are identified and encourage them to produce fixes or to reduce their impact**
- **Set up (private!) database of vulnerabilities**
 - To inform sites and developers
- **Urgent action → OSCT to manage**
- **After reaction time (45 days)**
 - vulnerability and risk analysis given to OSCT to define action – publication?
 - Will not publish vulnerabilities with no solution
- **Intend to report progress and statistics on vulnerabilities by middleware component and response of developers**
- **Balance between open responsible public disclosure and creating security issues with precipitous publication**

- **EGEE – OSG:**
 - Job submission demonstrated in both directions
 - Done in a sustainable manner
 - EGEE BDII and GIP deployed at OSG sites
 - Will also go into VDT
 - EGEE WN tools installed as a grid job on OSG nodes
 - Small fixes to job managers to set up environment correctly
- **EGEE – ARC:**
 - 2 workshops held (September, November) to agree strategy and tasks
 - Longer term want to agree standard interfaces to grid services
 - Short term:
 - EGEE→ARC: Try to use Condor component that talks to ARC CE
 - ARC→EGEE: discussions with EGEE WMS developers to understand where to interface
 - Default solution: NDGF acts as a gateway
- **In both cases:**
 - Catalogues are application choices – generally local catalogues use local grid implementations

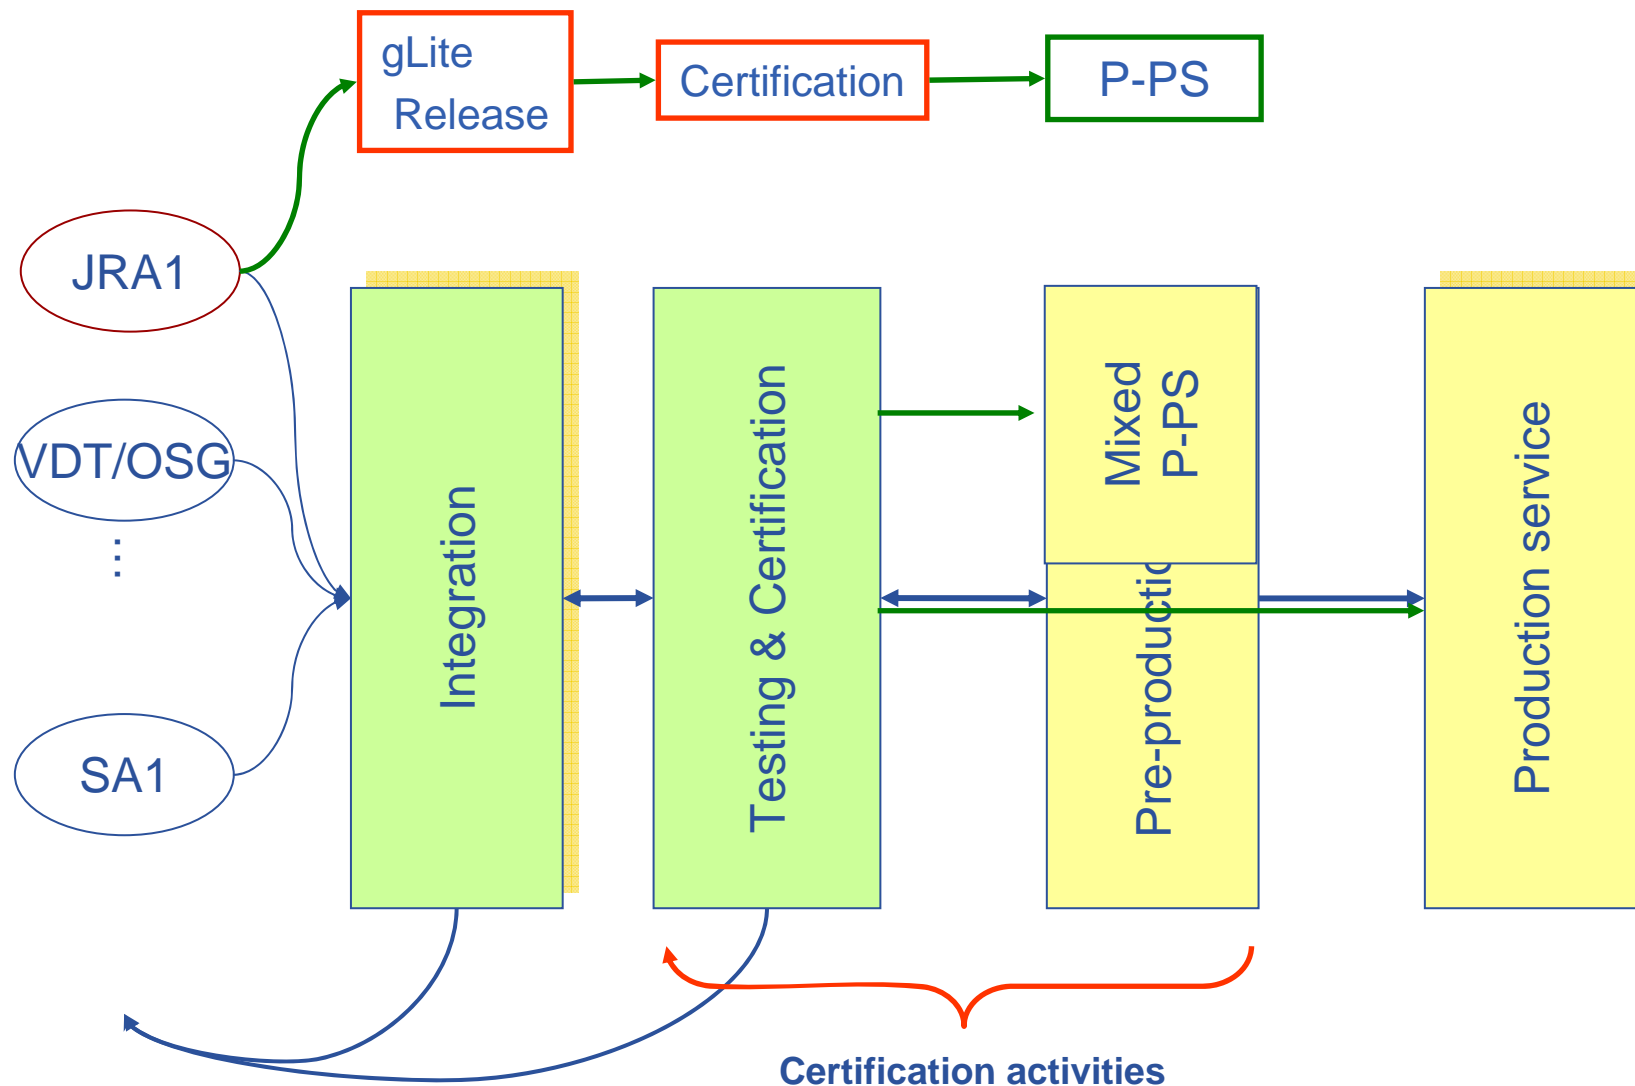
- **Goal: to improve level of “round-the-clock” operational coverage**
- **OSG have been to all of the EGEE operations workshops**
 - Latest was arranged as a joint workshop
- **Can we share operational oversight?**
 - Gain more coverage (2 shifts/day)
- **Share monitoring tools and experience**
 - Site Functional tests (SFT)
 - Common application environment tests
- **Strong interest from both sides**
- **User support workflows – interface**

- **Now: Write a short proposal of what we can do together**
 - Both EGEE and OSG have effort to work on this
- **Follow up in future operations workshops**

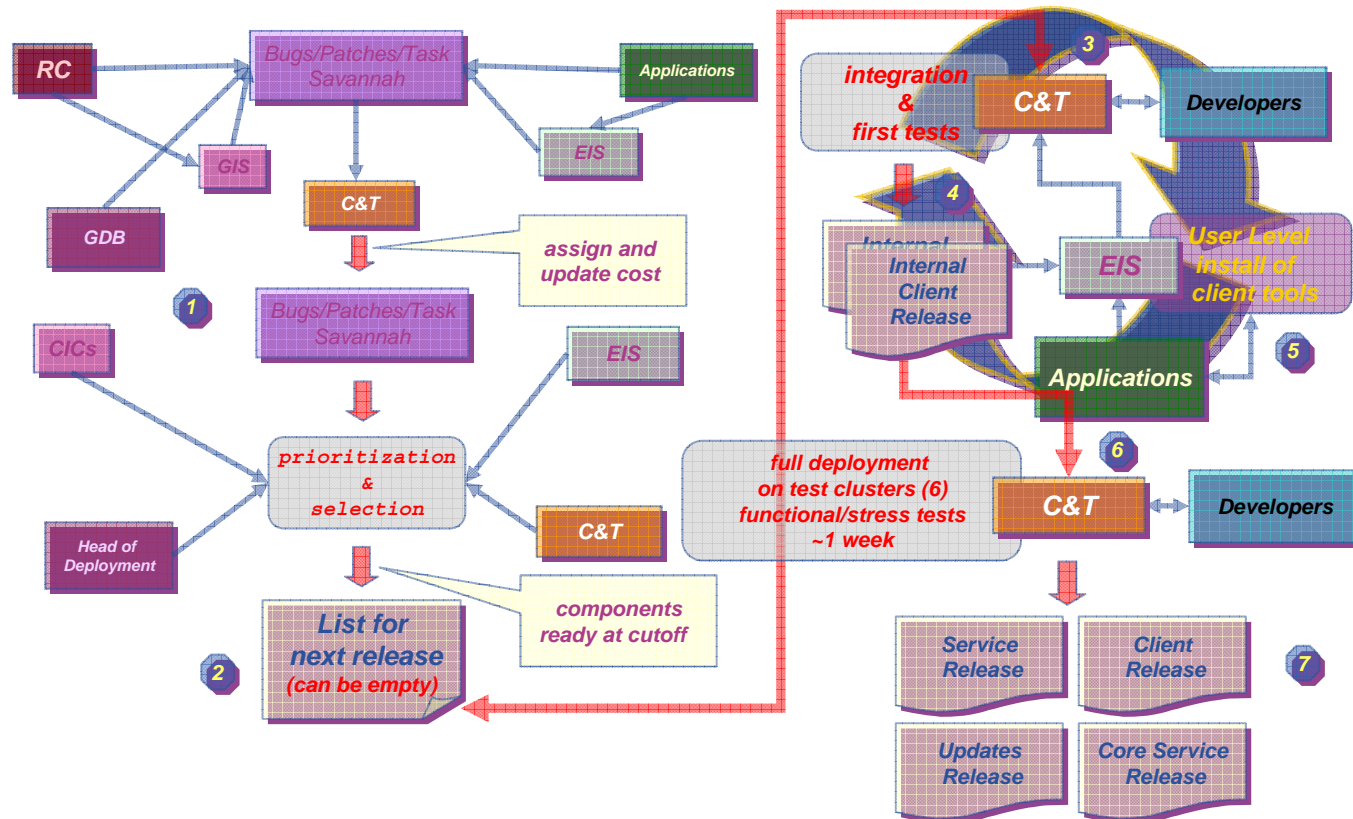
- **Interoperation and interoperability**
 - De-facto standards – common understandings/interfaces
 - GT2, GSI, SRM, BDII/GIP (MDS), ...
 - Agreement on schema:
 - GLUE 1.2/GLUE 2.0; GGF Usage record for accounting
 - *GLUE 2.0 will unify EGEE, OSG, ARC information schema*
 - Consider: common operations and job monitoring schema
- **Top-down vs bottom-up standards – must keep a balance in production**
 - What is working now (SRM, GLUE) vs what will help in future
 - Must maintain production service while introducing new components that apply standards → slow
- **Operations:**
 - SA1 “Cookbook”: summary of choices and experience deploying EGEE → intend to publish to GGF production grids
 - All aspects of operational security are very much collaborative with OSG and others (and very active in GGF)
 - Integration and certification is hard – standard interfaces and protocols should help
- **GGF Interoperability discussions**
 - Integrate bi-lateral interoperability work
 - EGEE/SA1 will contribute its work and experiences

- **Deployment process has improved significantly:**
 - Significant effort to improve the deployment process – better separation of functional improvements from critical updates
 - Simplified installation and configuration tool (YAIM) – made life much simpler for administrators
 - Wider deployment testing before release; also pre-production
 - GGUS coordinates problem follow up
- **Certification:**
 - Increased effort was identified (UK, INFN) to address lack of testing of new gLite components
 - Parallel processes to speed up gLite testing:
 - Production certification
 - “pure” gLite certification
 - Mixed (LCG-2.x + gLite) → this will become primary strategy


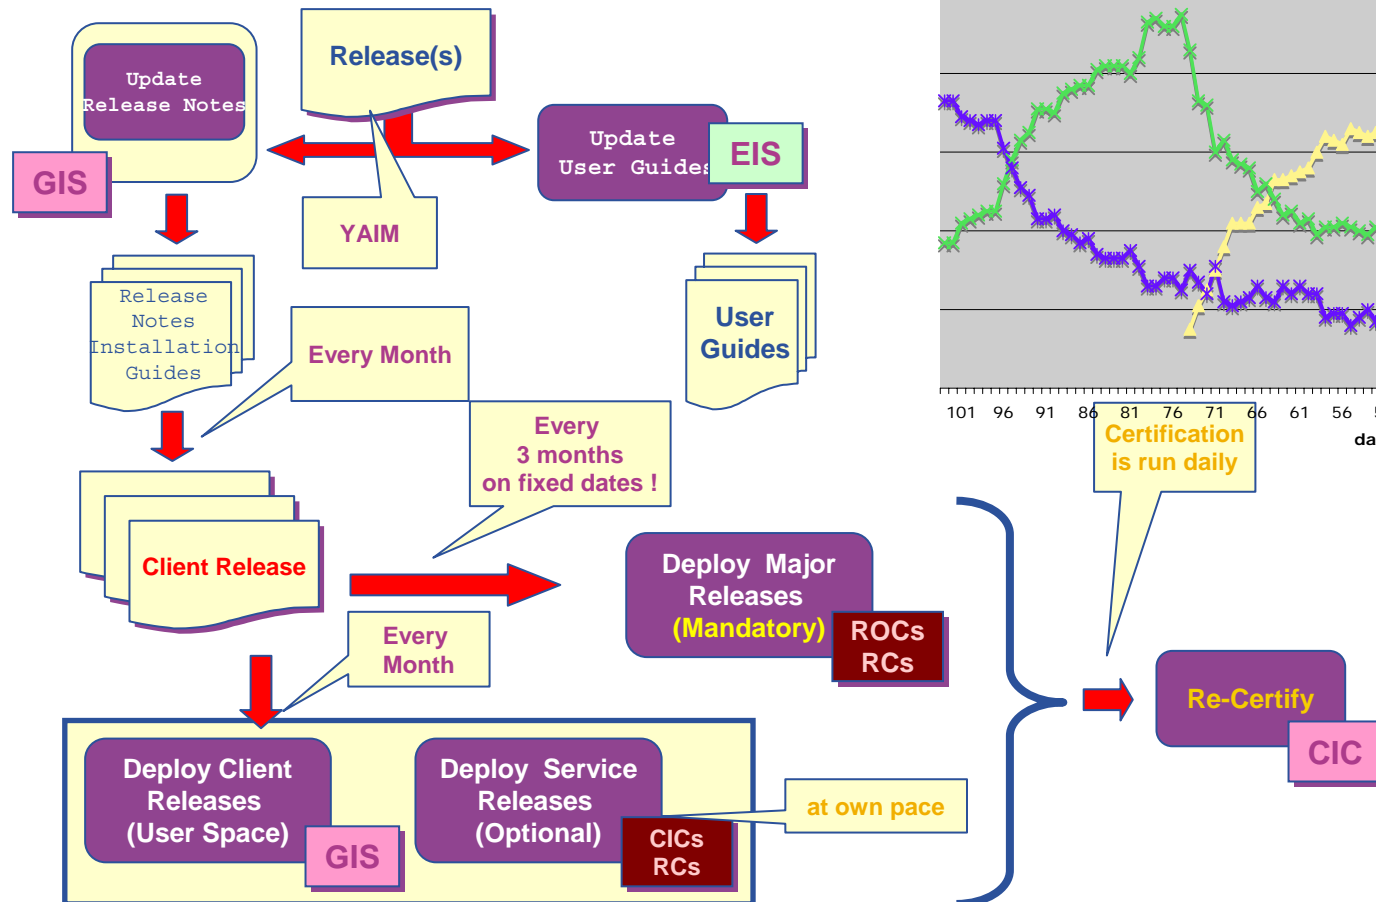
Moving components to production



- New process for major release was used (3 monthly fixed release)
 - Freeze of the candidate component list at a given date
 - Release at a given date (to allow planning)
 - 3 * 5 working days to upgrade the sites



- **Deployment ...**
 - Major releases have been expected to be installed within 3 weeks



- **Lessons learned**
 - Release definition non-trivial with 3 months intervals
 - Closing door for changes is almost impossible
 - Certification Tests need to be extended (performance tests)
 - Patches have to come with a standard set of information
 - Ports, variables, config changes....
 - Updates work quite well

- EGEE production service is a grid of **independent** federations
 - **ROCs schedule upgrades in their region**
 - 3 weeks upgrade window is neither realistic nor acceptable

 - Early announcement of new releases needed
 - **At -3 weeks**
 - *complete list of components and changes*
 - Problematic, because this means certification has to be finished
 - **At -2 weeks**
 - *deployment tests at: **ROC-IT, ROC-SE, ROC-UK***
 - **Last week** to implement feedback and final touches
 - Very useful – but cost of 3 weeks extra release time

- **Integrate JRA1 and SA1 processes**
 - Take into account these experiences gained over past 4 years
 - Ensure (TCG) priorities are driven by the applications

But:
 applications
 want rapid
 updates and
 deployment of
 new
 functionality

- **Finished certifying gLite 1.4.1**
 - Bugs found: 3 Critical ; 1 Normal ; 2 Minor
 - Will continue with more in depth testing.
 - Testbed also used to investigate bugs found elsewhere.
- **Certification testbed resides at CERN with 4 virtual sites.**
- **“Mixed” testbed (run and maintained by CNAF PPS)**
 - Contains both LCG-2 and gLite services
 - Carries out testing to compare LCG and gLite services (mostly RB so far)
 - Investigates the interoperability of gLite and LCG-2
- **The automated test suite is continually updated to cover new functionality provide by gLite.**

- **The current production middleware (“LCG-2”) is stable and is daily heavily used**
 - This has to be maintained as new components are added or components replaced
 - This will always be the case – there will always be new or better services coming
 - Thus, the production distribution must evolve in a controlled way that does not break existing applications but that adds new, or improves existing, functionality
- **There is a strong and reliable process in place**
 - Integration, testing, certification, pre-production, production
 - Process constantly evaluated and improved
 - All significant components of gLite 1.4/1.5 are either in production (R-GMA, VOMS, FTS) ...
 - ... or on the pre-production service (CE, WMS, Fireman, gliteIO)
 - Anticipate these being available in production distributions (alongside existing components at first) – by mid-2006 (many sooner)
- **The current LCG and gLite middleware will converge to a single distribution called gLite in early 2006**
- **Should not expect (or desire!) a big-bang switch to gLite (or anything else)**
- **Deploying in production any new software is a slow and time-consuming process, this lesson has been learned many times**

- **Accomplishments:**

- SA1 is operating world's largest grid infrastructure for science
- Significant resources available
- In use by many real production applications
 - 10K jobs/day
- Daily operations model is now well established
- User support process is in place and being used
 - But it is complex !
- Site stability is better controlled
 - Apps can select good sites
 - Understanding of metrics and what SLA might look like
- Ports to other architectures now exist
 - IA64, other Linuxes
- Convergence of middleware stacks under way
 - gLite components reaching production

- **Issues:**

- Hard to balance:
 - Needs of applications for rapid updates
 - Reliable scheduling wanted by sites
 - Adequate testing and certification
- Moving new middleware into production is time consuming:
 - Unrealistic expectations
 - Very stressful
 - But sw industry knows ...
- Essential to maintain stable production environment
 - While introducing new functionality, new services
 - Backwards compatibility
 - Expensive in resources and support
- Release of accounting (& other) data
 - some site policies restrict release of per-user data (privacy laws)
 - Accounting, job monitoring, ...
- Introducing new VOs is still too difficult

- **Remainder of EGEE**

- Milestones:

- MSA1.5 (PM21) – Expanded production grid available (50 sites)

- Deliverables:

- DSA1.7 (PM19) – Cookbook – internal review
- DSA1.8 (PM23) – Assessment of production operation (update of DSA1.4)
- DSA1.9 (PM21) – Release notes corresponding to MSA1.5

- Full metrics programme implemented (scope agreed in Pisa)

- Service availability SLA for LCG (MoU)

- Deploy major gLite components in production

- **Sustainability**

- Merge integration, testing (JRA1) with integration and certification (SA1) into single team with distributed partners

- Work with embryonic TCG to ensure application driven priorities reflected in development and deployment priorities

- Prepare processes for EGEE-II

- **EGEE/SA1 has deployed and is operating a production grid infrastructure many times the size of that anticipated**
 - 180 sites vs 50 as final milestone
- **The infrastructure is in daily use by many real applications from many disciplines**
 - 10K jobs per day, >1000 years cpu time in 2005
- **Daily operations is now an established (but constantly evolving procedure); stability is much better**
 - Metrics now established, better understanding of what is needed for SLAs
- **User support is complex, but process is in place and providing good service**
- **Convergence of production and gLite middleware stacks → gLite middleware distribution**