



Enabling Grids for
E-science in Europe

LCG/EGEE Grids & System Administration

Fotis Georgatos <gef@grnet.gr>
Grid Technologies Trainer, GRNET

University of Patras, December 7th-8th, 2006

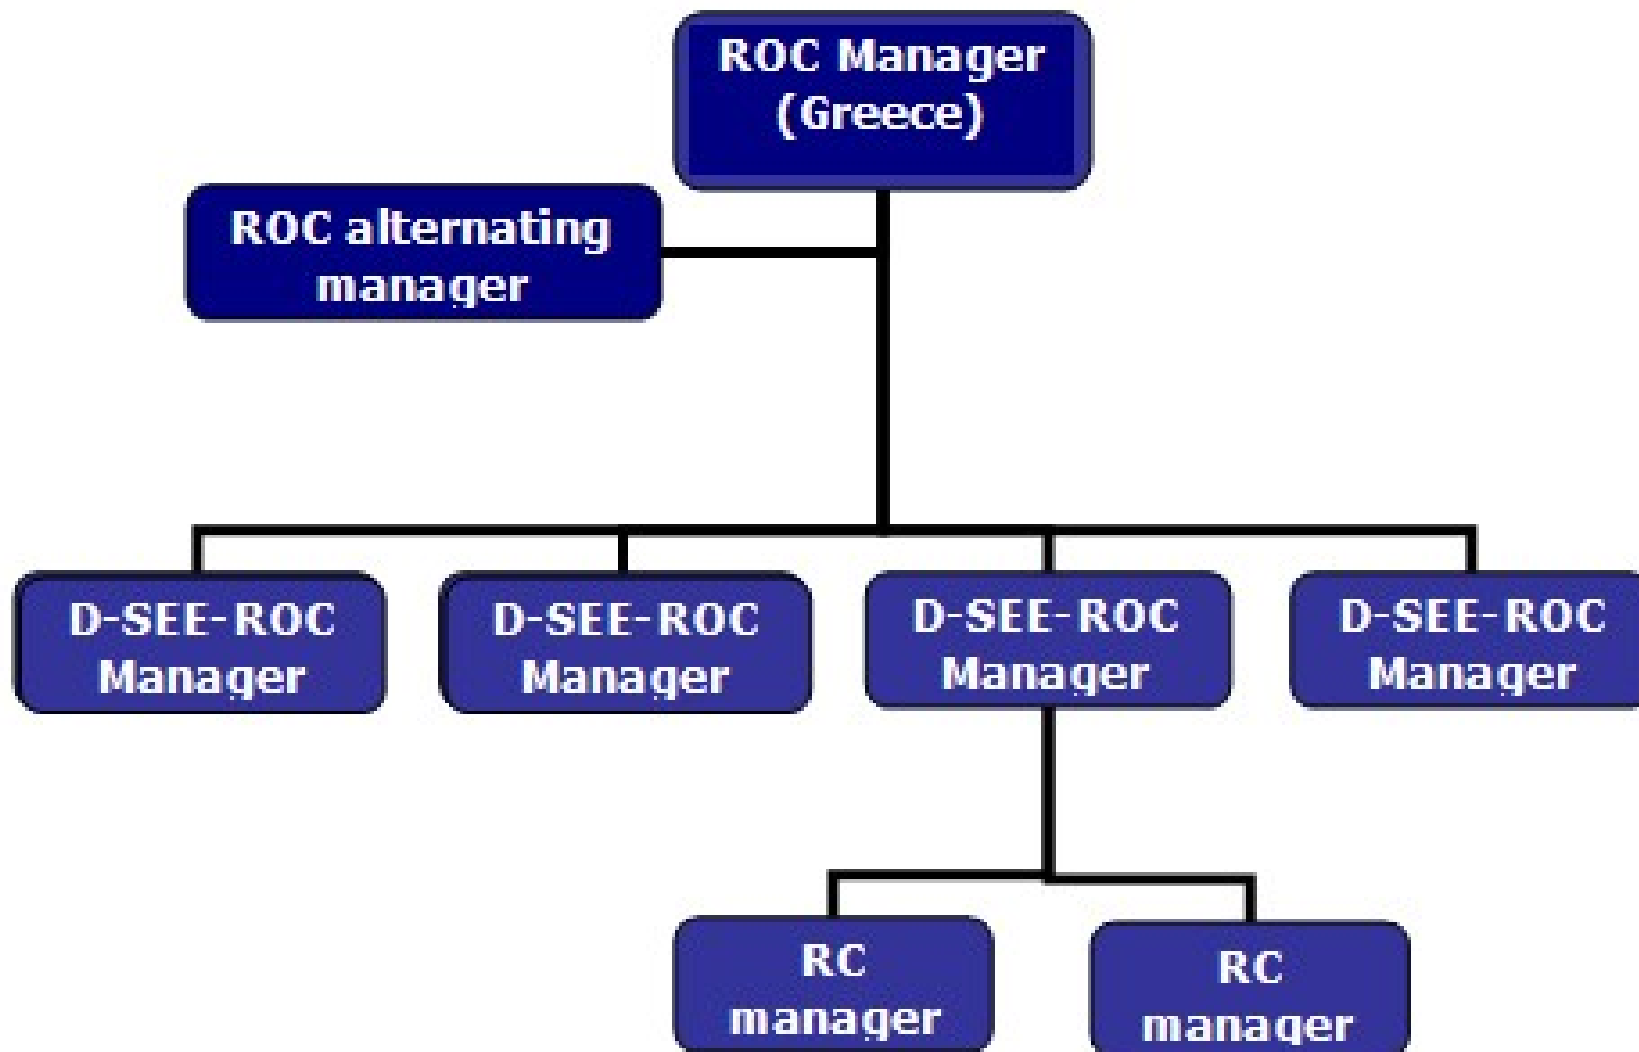


Watch out what you wish for!

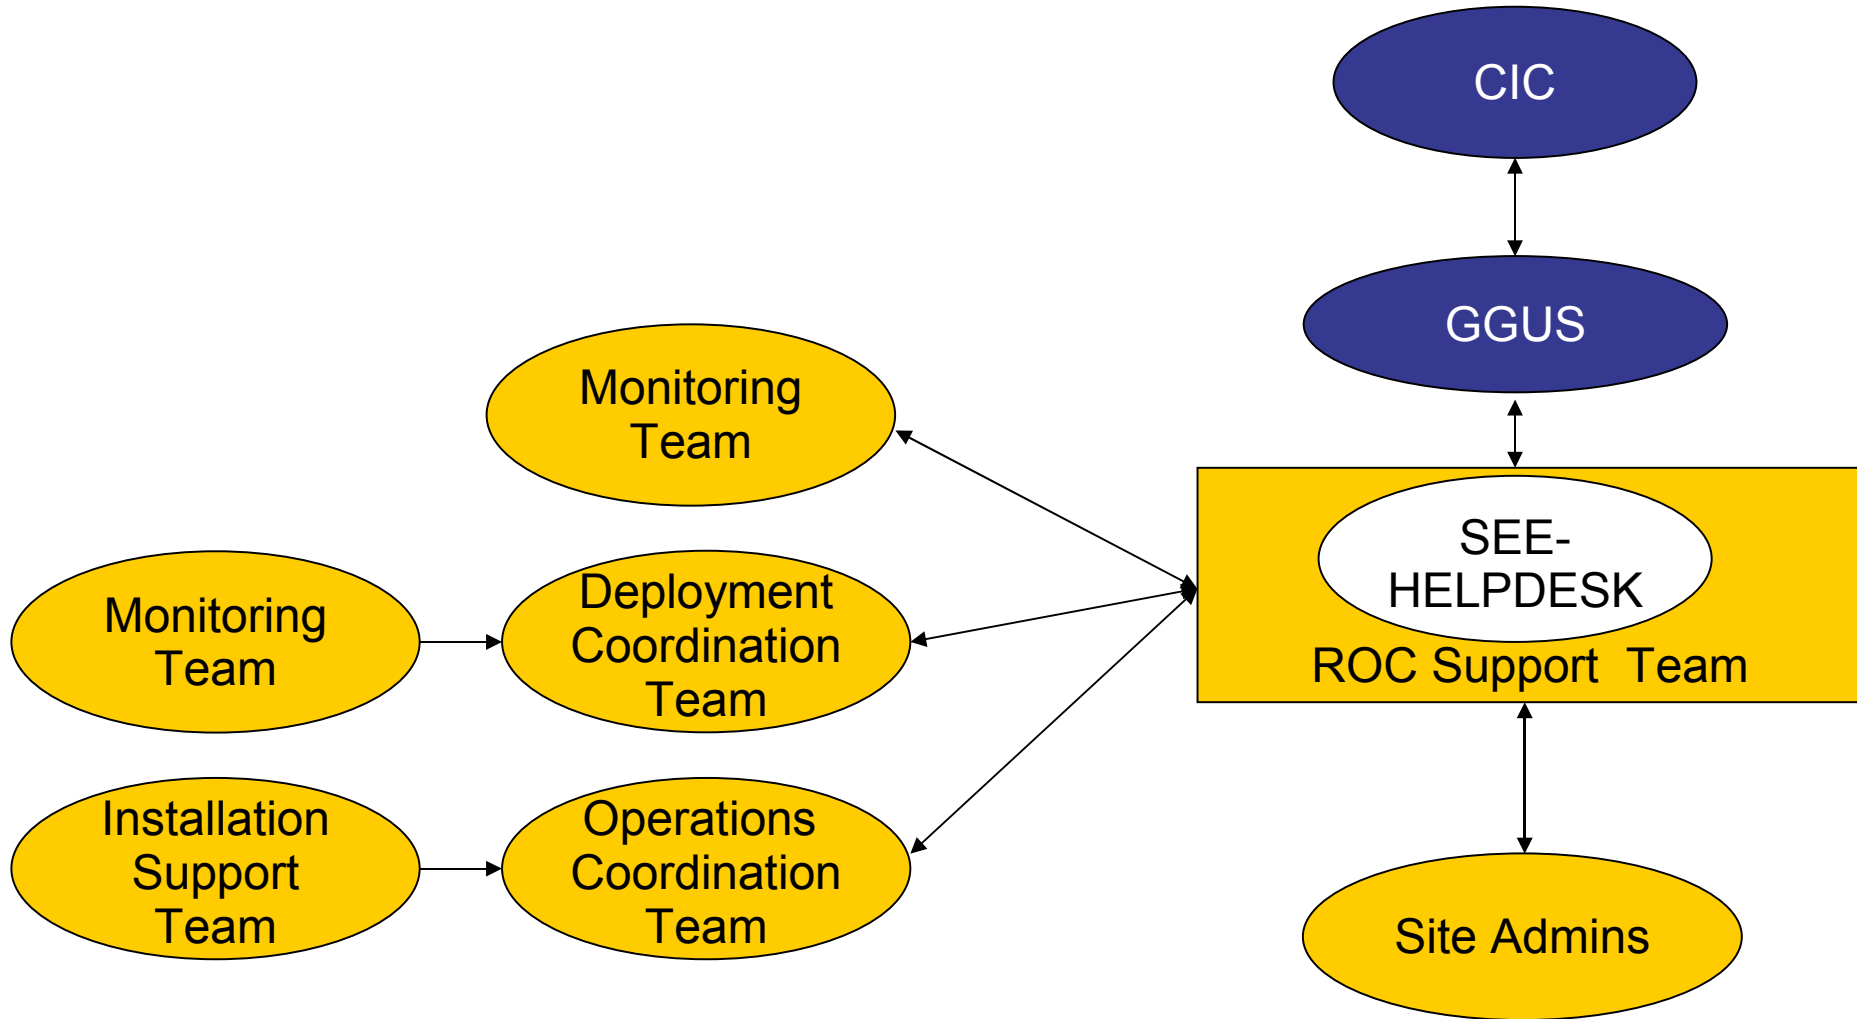


- Explain runtime operations and how the workflow will evolve
- Explain new site deployment and upgrade procedures
- Transfer experience gained during HellasGrid I & II phases
- What is GGUS, the egee-see helpdesk and other friends
- Runtime operations include
 - Site problem solving (ROC Support Team)
 - Minor upgrades (Operations Coordination Team)
 - Configuration changes (Operations Coordination Team)
 - Security incidents (Security Coordination Team)

The ROC structure within SEE



Prospective information workflow

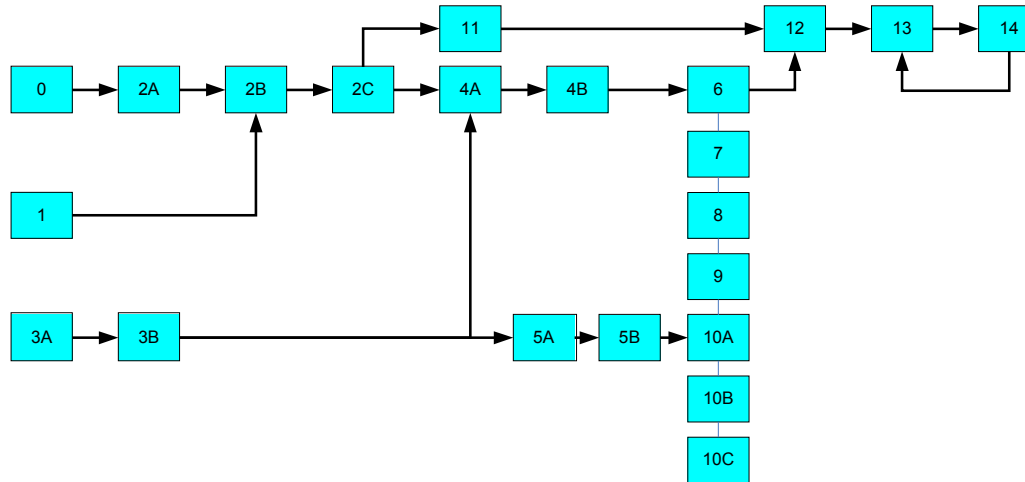


- Involves 3 steps:
 1. **Site registration**, done in coordination with the Deployment Coordination Team (DCT) – ‘candidate site’
 2. **Site installation**, done with the guidance and assistance of the Installation Support Team (IST) and the DCT - ‘uncertified site’
 3. **Site certification**, in coordination with the CIC-on-duty and the Site Certification Team (SCT) - ‘certified site’

Site induction procedure

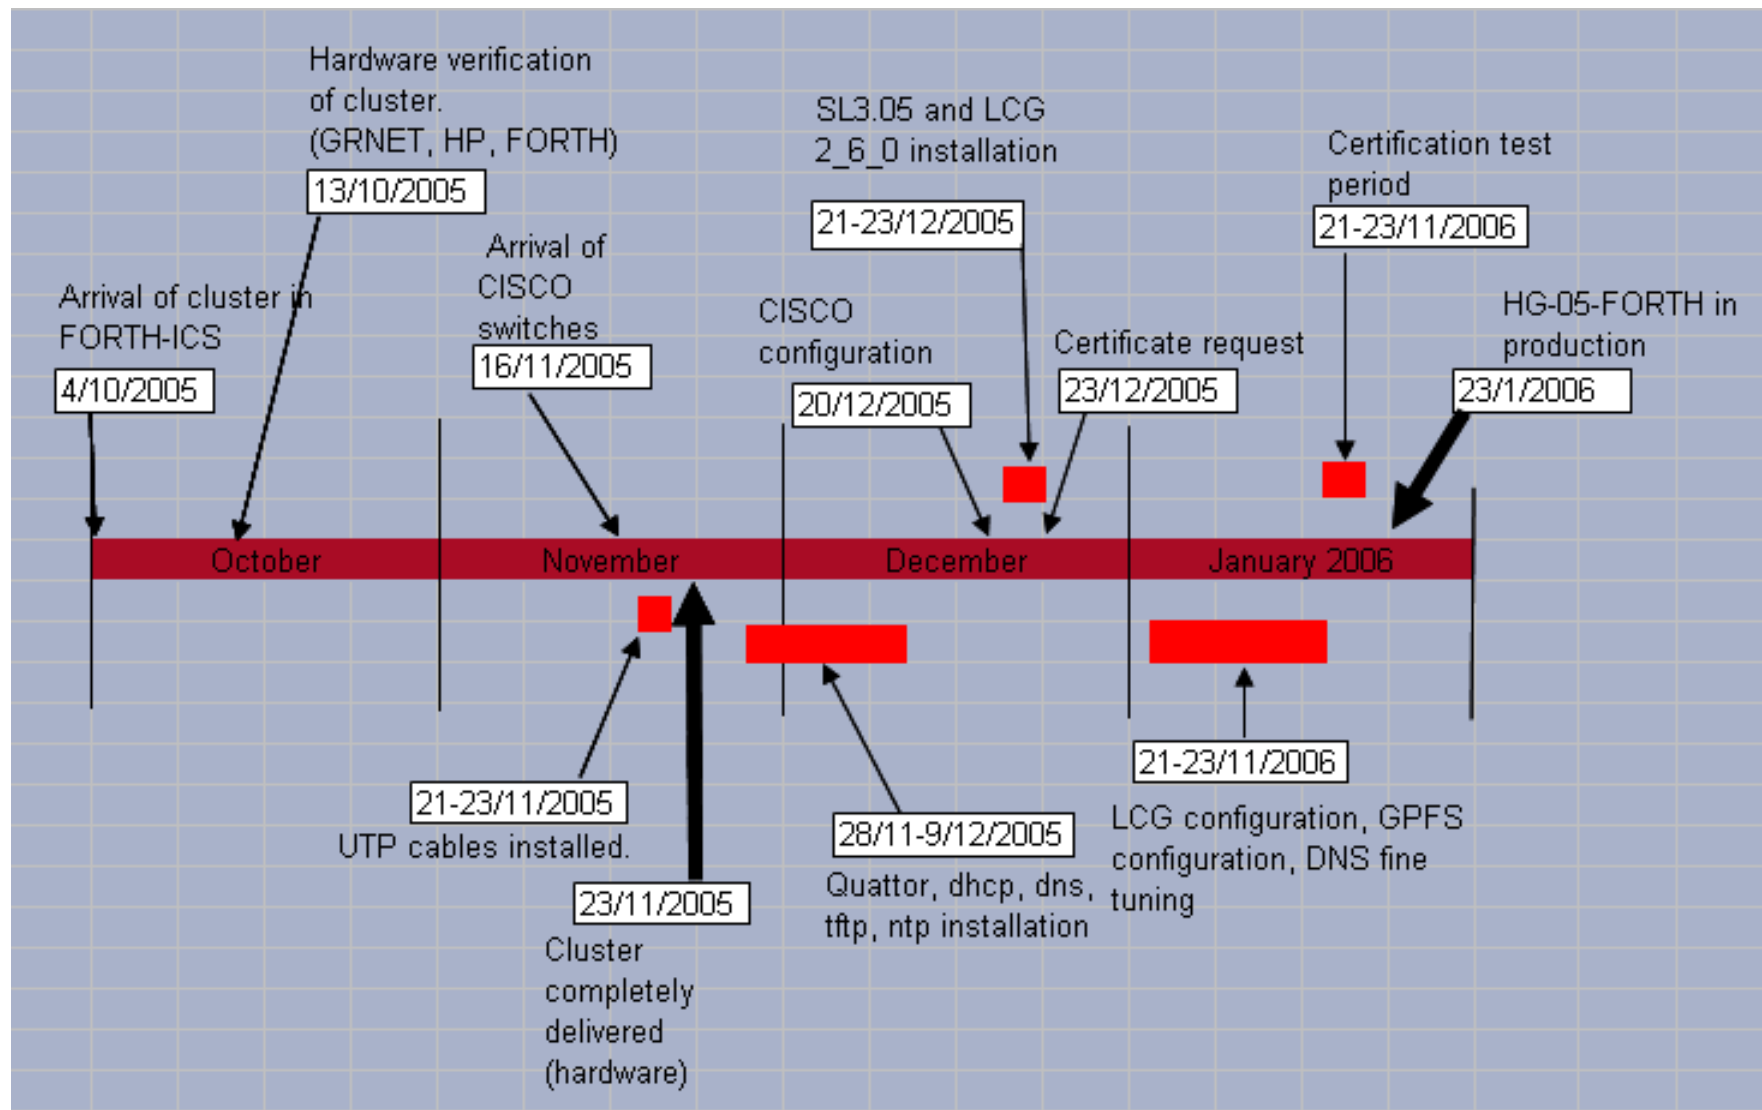
- Follow the 10-step registration procedure, as described here:
http://www.egee-see.org/Site_registration.php
- Perform a site installation according to the instructions here:
<http://lcg.web.cern.ch/LCG/Sites/releases.html>
- Follow the 5-step certification sequence:
http://www.egee-see.org/Site_certification.php
- Successful Site Functional Tests for five consecutive days, imply the site is stable enough to be considered “Production”
- DCT changes site status in the GOC database to ‘*certified*’
- From now on the site will submit the relevant weekly reports, also known as RC reports, Friday – Monday 11:00AM(GMT):
https://cic.in2p3.fr/index.php?id=rc&subid=rc_report&js_status=2

The complexities of a full site setup



- 0) UI account setup (Isabella or local)
- 1) An academic sends a fax to Hellasgrid CA, mentioning sysadmin's name
- 2A) User certificate is created by sysadmin(s), for himself
- 2B) Accept user certificate against CA, within a week's time!
- 2C) convert user certificate to pkcs12, add to browser and mail client
- 3A) Decide for IP, NTP and DNS configuration (addresses, names, servers)
- 3B) Verify NTP, DNS (forward and reverse) with host, www.dnsstuff.com etc
- 4A) Host certificates are created by sysadmin(s), (for CE, SE, MON ...)
- 4B) Accept host certificates against CA, within a week's time!
- 5A) Install SL30x to your systems
- 5B) Setup Java, NTP, lcg-yaim
- 6) Configure CE
- 7) Configure SE
- 8) Configure WNs
- 9) Configure MON 9B) & send its IP address to RGMA DB
- 10) Configure UI and/or 10B) BDII 10C) RB
- 11) Add site to GOCdb & inform ROC manager, update GOCdb
- 12) Ask Isabella's team for SFTs
- 13) Identify any erroneous entries with SFTs
- 14) Perform corrective actions for errors found in 13)

A real installation: HG-05-FORTH




Site registration procedure

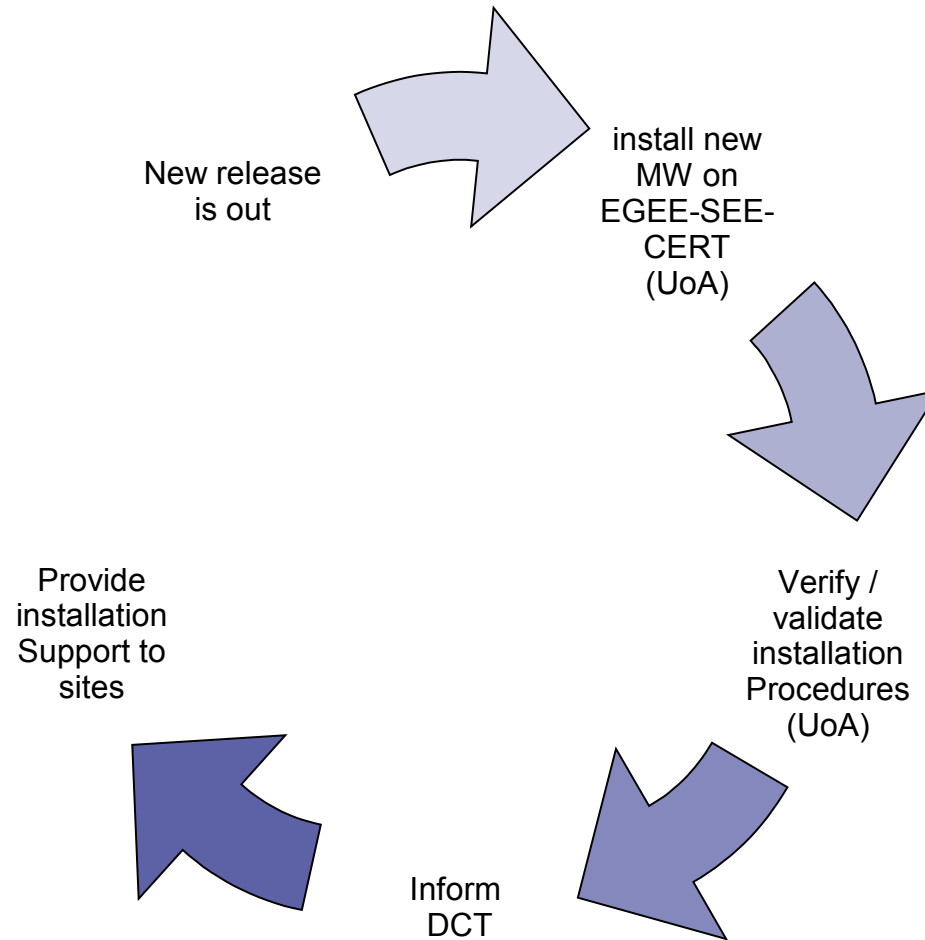
- Resource admin at candidate site obtains a certificate from Certification Authority (CA) eg. <http://ca.hellasgrid.gr> or <http://ca.egee-see.org>
- Site manager contacts the ROC Deployment Coordination Team (DCT)
Providing:
 2. Statement of acceptance of policy documents. <http://www.egee-see.org/aup>
 3. Site contact information
 4. HW details in the [Infrastructure.xls](#)
 5. Statement of agreement with **LCG/EGEE Security Incident Response procedures**
https://edms.cern.ch/file/428035/LAST_RELEASED/LCG_Incident_Response.pdf
- DCT validates the information and creates the new site's record in the GOC database; site status is '***candidate***'.
- Site to register with the helpdesk <https://helpdesk.egee-see.org/>
- Resource Admin at the site enters the remaining information in the GOC database, and then requests validation by the DCT
- The DCT validates the information and changes the site status to '***uncertified***'

- Done with the guidance and assistance of the IST and the DCT teams
 - Installation support through egee-see helpdesk
- Using the Installation Guidelines at
 - <http://lcg.web.cern.ch/LCG/Sites/releases.html>
- Supported OS: SL 30X, SLC30X, RHEL3 or similar


Site deployment/upgrade workflow



Deployment/upgrade workflow (Release Team)




Upgrade cases of the LCG Testbed



- **Time to upgrade**
~constant (~2.5 sites/day)
- **Takes a long time, $O(\#\text{sites})$, to upgrade entire infrastructure**
- **Better now than it was – site functional tests and operational oversight**
- **Need to move away from the need to do full upgrades more than 1-2 times / year**
 - But need to be able to deploy updates, new tools, security patches, etc.

Vicious circle of problem solving



Problem detection

- Done mostly by CIC-on-duty, but also ROC
- 1. Problem usually spotted using one of:
 - **GIIS Monitor:** <http://goc.grid.sinica.edu.tw/gstat/>
 - **Site Functional Tests:** <https://lcg-sft.cern.ch/sft/lastreport.cgi>
<https://lcg-sam.cern.ch:8443/sam/sam.py>
 - **GridICE:** <http://gridice2.cnaf.infn.it:50080/gridice/site/site.php>
 - **GOC Database:** <https://goc.grid-support.ac.uk/gridsite/db>
 - **GOC Site Map:** <http://goc02.grid-support.ac.uk/googlemaps/lcg.html>
 - **GOC Live Job Monitor:** <http://gridportal.hep.ph.ic.ac.uk/rtm/>
- 2. CICs open a Ticket at Savannah and/or GGUS (<https://savannah.cern.ch/projects/lcg2sites>) and sends an e-mail to site Administrators and the ROC **support team** grid-support@egee-see.org.
- **Site admins should look at:**
 - A) the GIIS Monitor and
 - B) the Site Functional Tests (or SAM) & start working on a problem

Monitoring tools: GISS Monitor

- GIIS Monitor (<http://goc.grid.sinica.edu.tw/gstat/>)

No	Site Reports	GIIS Host	bnode	cernse	gperf	sanity	serv	version	totalCPU	freeCPU	runJob	waitJob	seAvail TB	seUsed TB	maxCPU	avgCPU	DI	gice		
1	BG-INTRNE	ce1.inrne.bas.bg	ok	ok	ok	ok	ok	GLITE-3 0 2	26	8	18	0	0.90	0.00	27	24	OK	ok		
2	BG01-IPP	ce002.ipp.acad.bg	.	.	ok	warn	ok	na	18	9	9	154	0.90	0.10	18	15	OK	OK	info	
3	BG02-IM	ce001.imbm.bas.bg	.	.	ok	ok	ok	GLITE-3 0 2	3	2	1	0	0.02	0.03	3	2	OK	ok		
4	BG04-ACAD	ce02.grid.acad.bg	.	.	ok	ok	ok	GLITE-3 0 2	80	60	20	0	0.04	0.01	80	78	OK	ok		
5	BG05-SUGrid	ce001.grid.uni-sofia.bg	.	.	ok	ok	ok	GLITE-3 0 0	14	8	6	8	0.03	0.04	14	4	OK	ok		
6	CY-01-KIMON	ce101.grid.ucy.ac.cy	ok	ok	ok	ok	ok	GLITE-3 0 1	70	51	19	0	0.19	0.01	72	69	OK	ok		
7	GR-01-AUTH	node001.grid.auth.gr	.	.	ok	ok	ok	GLITE-3 0 0	14	11	3	9	0.17	0.03	14	12	OK	.		
8	GR-03-HEPNTUA	ce.hep.ntua.gr	.	.	ok	ok	ok	GLITE-3 0 2	22	18	10	0	0.78	0.06	49	23	OK	ok		
9	GR-04-FORTH-ICS	grid001.ics.forth.gr	.	.	ok	note	ok	GLITE-3 0 0	5	2	3	36	0.00	0.01	5	4	OK	ok		
10	GR-05-DEMOKRITOS	ikaros4.inp.demokritos.gr	.	.	ok	ok	ok	GLITE-3 0 0	24	18	0	0	1.68	0.00	44	40	SD	.		
11	GR-06-IASA	ce02.marie.hellasgrid.gr	.	.	ok	ok	ok	GLITE-3 0 1	20	2	18	1	0.17	0.00	20	17	OK	info		
12	HG-01-GRNET	ce01.isabella.gnet.gr	ok	ok	ok	ok	ok	GLITE-3 0 2	64	1	63	7	3.32	1.46	64	63	OK	ok		
13	HG-02-IASA	ce01.marie.hellasgrid.gr	.	.	ok	ok	ok	GLITE-3 0 1	118	40	78	0	2.68	0.17	118	117	OK	info		
14	HG-04-CTI-CEID	ce01.kallisto.hellasgrid.gr	.	.	ok	ok	ok	GLITE-3 0 0	120	30	93	0	2.47	0.15	120	116	OK	ok		
15	HG-05-FORTH	ce01.ariagni.hellasgrid.gr	.	.	ok	ok	ok	GLITE-3 0 0	116	2	116	9	2.69	0.22	120	99	OK	ok		
16	HG-06-EKT	ce01.athena.hellasgrid.gr	ok	ok	ok	ok	ok	GLITE-3 0 2	226	1	225	17	9.37	0.00	228	207	OK	OK	ok	
17	LCG-IL-OU	grid01.cslab.openu.ac.il	.	.	ok	ok	ok	GLITE-3 0 0	10	1	9	27	0.11	0.01	10	9	OK	ok		
18	TAU-LCG2	lcfgng.cs.tau.ac.il	.	.	ok	ok	ok	GLITE-3 0 1	19	2	30	33	0.08	0.01	22	11	OK	.		
19	TECHNION-LCG2	ds-lcg-ce01.cs.technion.ac.il	.	.	ok	ok	ok	GLITE-3 0 2	13	9	0	0	0.09	0	18	3	JL	.		
20	WEIZMANN-LCG2	wipp-ce.weizmann.ac.il	ok	ok	ok	ok	ok	GLITE-3 0 1	50	20	26	8	0.09	0.06	50	48	OK	ok		
21	MK-01-UKIM_II	grid-ce.ii.edu.mk	.	.	ok	ok	ok	GLITE-3 0 0	12	9	3	0	0.21	0.00	12	11	OK	ok		
22	NIHAM	alice003.nipne.ro	.	.	ok	ok	ok	GLITE-3 0 1	2	2	0	0	0.04	0.00	2	1	OK	ok		
23	RO-01-ICI	testbed001.grid.ici.ro	.	.	ok	ok	ok	GLITE-3 0 1	20	16	2	0	0.40	0.01	20	19	OK	.		
24	RO-02-NIPNE	tbat01.nipne.ro	.	.	ok	ok	ok	GLITE-3 0 1	102	86	16	0	0.36	0.05	102	101	OK	ok		
25	TR-01-ULAKBIM	ce.ulakbim.gov.tr	ok	ok	ok	ok	ok	GLITE-3 0 2	85	54	12	10	2.36	0.69	108	95	JS	JS	ok	
26	AEGIS01-PHY-SCL	ce.phy.bg.ac.yu	ok	ok	ok	ok	ok	GLITE-3 0 2	101	1	100	200	0.05	0.08	101	97	OK	OK	ok	
									sites	countries		totalCPU	freeCPU	runJob	waitJob	seAvail TB	seUsed TB	maxCPU	avgCPU	
									Total	26	8	1354	463	880	519	29.32	3.33	1441	1285	

Monitoring tools: SFTs

- Site Functional Tests (<https://lcg-sft.cern.ch/sft/lastreport.cgi>)

Test summary

	SD	JL	JS	CT	OK	total
dteam			1		5	6

Colours definition

SD	Scheduled downtime	#a3a3a3
JL	Job list match failed	#aab3ff
JS	Job submission failed	#f4876b
CT	Critical tests failed	#f9d48e
NT	Non-critical tests failed	#f2f98e
OK	OK	#b2f98e

Test abbreviations

csh	CSH test
rgmasc	R-GMA Secure Connector
swdir	VO software directory
rgma	R-GMA
wn	WN host name
ver	Software Version (WN)
ca	CA certs version
rm	Replica Management
votag	VO Tag management
js	Job submission
bi	BrokerInfo
apel	Apel test

	St.	Region	Site Name	Site CE	VO dteam													
					St.	js	wn	ver	ca	rgma	bi	csh	rm	votag	swdir	rgmasc	apel	
1.	OK	SouthEasternEurope	BG01-IPP	ce001.grid.bas.bg	OK	O	I	2	7	0	O	O	O	O	O	O	O	O
2.	JS	SouthEasternEurope	EGEE-SEE-CERT	ce01.gridctb.uoa.gr	JS	X	??	??	??	??	??	??	??	??	??	??	??	??
3.	OK	SouthEasternEurope	GR-01-AUTH	node001.grid.auth.gr	OK	O	I	2	7	0	O	O	O	O	O	O	O	O
4.	OK	SouthEasternEurope	HG-01-GRNET	ce01.isabella.grnet.gr	OK	O	I	2	7	0	O	O	O	O	O	O	O	O
5.	OK	SouthEasternEurope	HG-02-IASA	ce01.marie.hellasgrid.gr	OK	O	I	2	7	0	O	O	O	O	W	O	O	O
6.	OK	SouthEasternEurope	NIHAM	alice003.nipne.ro	OK	O	I	2	7	0	O	O	O	O	W	O	O	O

The improved SFT tool is called SAM: <https://lcg-sam.cern.ch:8443/sam/sam.py>

Monitoring tools: GridICE

- Grid ICE (<http://gridice2.cnaf.infn.it:50080/gridice/site/site.php>)

The screenshot shows the GridICE monitoring interface. At the top, there are navigation tabs: Geo view, Site view, VO view, Help, and About. Below these, there are sub-tabs: General, Gris, Host, Job, Charts, and Network. The main content area displays a table of site resources, categorized into Computing Resources and Storage Resources. The table includes columns for Site, Region, GK#, Q#, RunJob, WaitJob, JobLoad, Power, WN#, CPU#, CPUload, Available, Total, %, and MH#.

Site ▼	Region	GK#	Q#	RunJob	WaitJob	Computing Resources					Storage Resources			MH#
						JobLoad	Power	WN#	CPU#	CPUload	Available	Total	%	
AEGIS01-PHY-SCL	SEE	1	8	100	198	100%	0	25	97	100%	-	-	-	31
BG01-IPP	SEE	2	22	18	306	-	-	-	-	-	33.5 GB	105.1 GB	68%	-
BG02-IM	SEE	1	6	1	0	5%	0	3	2	5%	13.6 GB	32.9 GB	59%	2
BG04-ACAD	SEE	-	-	-	-	83%	0	40	80	23%	-	-	-	43
BG05-SUGrid	SEE	1	8	4	16	5%	0	4	10	0%	38.5 GB	83.5 GB	54%	2
CY-01-KIMON	SEE	1	9	26	0	0%	0	36	70	17%	-	-	-	4
GR-01-AUTH	SEE	1	11	5	9	-	-	-	-	-	180.3 GB	217.6 GB	17%	-
GR-03-HEPNTUA	SEE	-	-	-	-	77%	0	11	22	56%	-	-	-	15
GR-04-FORTH-ICS	SEE	-	-	-	-	100%	0	4	3	100%	-	-	-	7
GR-05-DEMOKRITOS	SEE	-	-	-	-	95%	15K	10	40	98%	864.6 GB	868.5 GB	0%	4
HG-01-GRNET	SEE	-	-	-	-	100%	0	23	61	96%	-	-	-	28
HG-02-IASA	SEE	-	-	-	-	-	-	-	-	-	-	-	-	1
HG-04-CTI-CEID	SEE	1	14	109	0	100%	0	60	120	63%	2.4 TB	2.5 TB	6%	63
HG-05-FORTH	SEE	1	14	116	22	100%	0	58	114	91%	-	-	-	63
HG-06-EKT	SEE	1	15	225	18	100%	0	113	224	91%	-	-	-	118

- There can also exist local MON instances (eg: <http://mon.egee-see.org>)

.. and many other monitoring tools...

The screenshot displays a monitoring interface with three main components:

- Google Earth:** Shows a satellite view of a city area with a red pin labeled "GR-05-DEMOKRITOS". An inset map of Europe shows various colored markers representing different locations.
- Application Traffic Breakdown:** A line chart titled "Application Traffic Breakdown: University of Crete" showing inbound and outbound traffic in Mbit/s from 13:30 to 14:20. The chart is color-coded by application. Below the chart is a summary table:

	In	Out
Total	37.48 bps	-117.77 bps
BitTorrent	19.01 bps	-60.62 bps
eDonkey	8.47 bps	-49.95 bps
DC++	0.67 bps	-1.14 bps
Gnutella	1.53 bps	-0.11 bps
FTP	0.06 bps	-0.00 bps
HTTP	6.12 bps	-3.26 bps

- Private Section Non-Anonymized:** Two tables showing traffic statistics for specific IP addresses.

INCOMING TRAFFIC		
IP	Protocol	Traffic
147.52.3.XXX	Gnutella	1007.58 Kbps
147.52.3.X	BitTorrent	602.90 Kbps
147.52.99.XXX	DC++	554.04 Kbps
147.52.67.XXX	HTTP	484.45 Kbps
147.52.19.XXX	BitTorrent	407.29 Kbps
147.52.48.XXX	BitTorrent	401.82 Kbps
147.52.19.XX	BitTorrent	388.10 Kbps
147.52.99.X	BitTorrent	350.39 Kbps
147.52.99.X	eDonkey	316.06 Kbps
147.52.19.XXX	BitTorrent	273.71 Kbps

OUTGOING TRAFFIC		
IP	Protocol	Traffic
147.52.99.X	eDonkey	3952.56 Kbps
147.52.3.X	BitTorrent	3317.29 Kbps
147.52.19.XXX	BitTorrent	3204.69 Kbps
147.52.99.XX	eDonkey	2955.83 Kbps
147.52.19.XX	eDonkey	1795.67 Kbps
147.52.48.XXX	BitTorrent	1718.36 Kbps
147.52.99.XXX	DC++	1584.79 Kbps
147.52.19.XX	BitTorrent	969.58 Kbps
147.52.99.X	BitTorrent	937.98 Kbps
147.52.19.XXX	BitTorrent	672.27 Kbps

TOTAL TRAFFIC			
INCOMING TRAFFIC		OUTGOING TRAFFIC	
IP	Traffic	IP	Traffic
147.52.99.X	706.91 Kbps	147.52.99.X	5196.19 Kbps

- Read the description of the task / problem
 - mail received usually from the CIC-on-duty or ROC: helpdesk.egee-see.org
- Look (**yes!**) at “Sites Functional Tests”/SAM for more details
- Get assistance / guidance from
 - Knowledge databases with information about the typical problems
 - <http://goc.grid.sinica.edu.tw/gocwiki/SiteProblemsFollowUpFaq>
 - <http://grid-it.cnaf.infn.it/index.php?knowledgebase>
 - <http://www.gridpp.ac.uk/tb-support/faq/index.html>
 - The pool of expertise that resides in mailing lists:
 - **LCG-ROLLOUT@LISTSERV.RL.AC.UK**
 - egee-sa1-tech@grnet.gr

- Try to **reproduce** the problem and **identify** its cause
- Try to **solve** the problem
- If needed request assistance from egee-sa1-tech@grnet.gr and/or the CIC-on-duty
- If problem is solved
 - Report to **Savannah/GGUS** that it is solved, explain briefly what caused it
 - CIC-on-duty or the ROC support team updates the ticket to “Site OK”
 - Your site will be in quarantine for 3 working days
- **Else**
 - Report to **Savannah/GGUS** and grid-support@egee-see.org the reason why the problem is not solved.
 - Keep on trying ... 😊


Where to get information from

- ICCS team Grid Certification Guide (you should be familiar!)
http://www.egee-see.org/content/modules/downloads/Certification_v2.pdf
- HOW TO TEST AN LCG2 SITE (LCG2-Site-Testing.pdf)
<http://grid-deployment.web.cern.ch/grid-deployment/documentation/LCG2-Site-Testing.pdf>
- DESY's Test Guide (find examples and debug commands)
<http://grid.desy.de/tests/>
- GridPP: FAQ for LCG Site Administrators
<http://www.gridpp.ac.uk/tb-support/faq/>
- GOCwiki: Grid Administration FAQs and Troubleshooting
<http://goc.grid.sinica.edu.tw/gocwiki>
- Information System Troubleshooting (GRISs, GIISs, BDIIs)
<http://lfield.home.cern.ch/lfield/trouble.html> (if you have to debug the MDS)


- Operational since March 2005
- Purpose:
act as catch-all VO for new applications of Regional Users
- How to enroll : Get a certificate and go to
<https://www.grid.auth.gr/services/voms/SEE/request.php>
- Evaluation committee for new user/applications (SEE-EGAAP).
Formal evaluation procedure will be established based on
EGAAP existing policies but with relaxed requirements.

Hierarchical National Infrastructure


**Tier 0 –
Demokritos Located,
Central GRNET node
(64 CPUs-10TB SAN)**



**Tier 1 –
5 peripheral nodes:
Athens (2), Patras,
Thessaloniki, Crete**

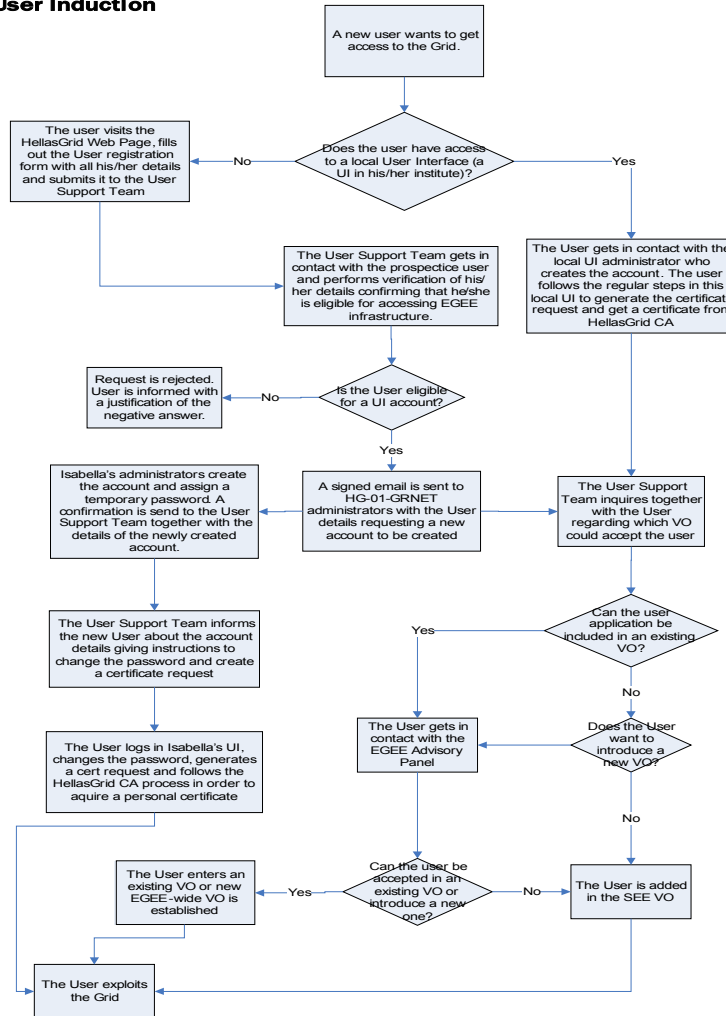


**Tier 2 – Integration of
other infrastructure
(sch.gr, Laboratories)**




Policies and Procedures

New User Induction



- Policies exist for nearly any grid related activity, such as:
 - Adding a site (ROC-related)
 - Having a user get a certificate
 - Adding a user to a VO (AUP)
 - Adding a system administrator
 - Handling a security incident
 - ...more we don't know...
- LCG Grid evolves faster than the policies themselves, so do “handle with care” 😊


Ανατομία ενός VO: SEE VO



- Κατάλογος χρηστών:
 - VO server & Myproxy
- Κατάλογος πόρων:
 - BDII (LDAP based!)
- Υπολογιστικοί πόροι:
 - Resource Broker (RB)
- Χωρητικοί πόροι:
 - LCG File Catalog (LFC)
- Τοπικές Υποδομές:
 - CE & WNs, SE, UI κλπ.

- Various channels have been established for User Support
- Problem Reporting
 - EGEE-SEE Helpdesk (<https://helpdesk.egee-see.org/>)
 - GGUS (<https://gus.fzk.de>) EGEE-wide
- Localized Documentation, FAQs, Tips, Instructions
 - EGEE-SEE Wiki (<http://wiki.egee-see.org/index.php/Users>)
 - EGEE-SEE Web Site (<http://www.egee-see.org/>). Recently restructured to provide easier navigation for the users.
- Other ad-hoc resources (various web sites, documents etc.)

Υποδομή HellasGrid I, Isabella




- High Availability
 - Through HW and SW redundancy
- Security aspects thoroughly examined
- Timely Resolution of problems
 - Efficient collaboration between team members
 - Close cooperation with VOs
- Aid in the deployment of New Sites / certification

- Redundant disks on Service Nodes and WNs (RAID1 - mirroring)
- Redundant Storage Infrastructure for SE/SAN at multiple levels
 - RAID5 volumes on storage array
 - Redundant FC disk controllers + PSUs.
 - Redundant FC links in **fail-over** mode AND **balancing** mode for GPFS storage nodes
 - Node redundancy at the GPFS level

- Redundant GPFS storage nodes
 - One primary / one secondary per Network Storage Device (NSD)
- Redundant network service instances
 - For DNS two on-site, two off-site servers
 - Similar redundancy in handling the NTP protocol

Security: OpenVPN



- Management interfaces unreachable from the outside
- Secure remote access to management VLAN using the free OpenVPN tool
 - Certificate-based authentication
 - SSL-based encryption

- Nodes are kept in a security hierarchy with different levels
 - Platinum: Backup server, Remote Console Access
 - Gold: Management Server, HW monitoring
 - Copper: Services & Worker Nodes for the Grid
- Explicitly defined trust relationships!

- System-based Intrusion Detection Systems
 - OSS tools (e.g. chkrootkit)
 - Custom-written scripts (chksetuid with md5sum)
- Network based Intrusion Detection Systems
 - Snort
- *Factorized* Logging infrastructure based on syslog-ng
 - Logs collected centrally at the management server
 - Logs replicated at off-site location

- **Support contracts with vendors in place**
 - IBM is the provider of the hardware of the HellasGrid I node
 - HP is the provider of the hardware of the HellasGrid II nodes
 - Support contracts are in place, in order to ensure timely replacement of components that are found to malfunction. It has been used in practice, equipment needs maintenance after all.
 - Support contracts extend to equipment such as the Uninterruptible Power Supply and the Air-Conditioners
 - Vendors are willing to keep up their promises, but the complexity of the equipment and its configuration setup should not be underestimated

Day-to-day Operations: Shifts

- Two shifts per day
 - 09:00 – 15:00, 15:00-21:00, Monday to Friday
- Ensures proper response in case of an emergency
- Handling of tickets coming from EGEE & SEE
- Timely service of requests from User Support team

- EGEE Site Functional Tests
- Local monitoring tools
 - Ganglia
 - MRTG
- Vendor-specific tools
 - IBM Cluster Systems Management
 - Monitors various node health parameters
 - Sends e-mail alerts which can be routed to mobiles

- Streamlining of new site installations
 - Guide for new HW installations
 - Customized instructions for OS deployment
- Certification Period
 - Certification SFTs run by the HG-01-GRNET team for all yet uncertified sites
 - Site enters production when the tests have run without problems for 5 days

Collaboration of Team Members

- **Request Tracker**
 - Web-based Ticketing System
 - Used for day-to-day collaboration
 - Permanent archive of information on all events during shifts
 - Facilitates integration of new team members
 - Knowledge base for all of HellasGrid' clusters
- **Weekly meetings**
 - Review of open tickets
 - Planning of future activities

RT ticketing system: the big picture

The screenshot shows the RT ticketing system interface in Mozilla Firefox. The browser window title is "RT at a glance - Mozilla Firefox". The address bar shows "https://rt.cslab.ece.ntua.gr/". The page title is "CSLab" and the URL is "RT for cslab.ntua.gr". The user is logged in as "gef".

The interface displays "RT at a glance" with a search bar and a dropdown menu for "New ticket in" set to "egee-3rd-parties". The main content area is divided into two sections:

- 10 highest priority tickets I own...**
- 10 highest priority tickets I requested...**

There is also a "Quick search" table on the right side of the page.

# Subject	Queue	Status	Owner
4601 test greek iso8859-7	General	open	
5651 hwclock: erroneous delta times?	General	open	

# Subject	Queue	Status	Owner
810 Να βρεθεί υποστήριξη και upgrades για airco/ups κλπ.	HG-01-GRNET	open	kkoum
5651 hwclock: erroneous delta times?	General	open	gef

Queue	New	Open
egee-3rd-parties	0	0
egee-sa1-tech	0	0
General	1	6
grid-sec	0	21
HG-01-GRNET	8	16
HG-01-GRNET-LOGS	0	0
HG-06-EKT	1	1
HG-06-EKT-CSIRT	1	0
hwinfo	0	0
lcg-rollout-digest	286	234
root_director	843	106
smokealert	0	24
test	0	0
vo-sites	886	364

The bottom of the browser window shows the Windows taskbar with the Start button, taskbar icons for Firefox, Outlook, and PowerPoint, and the system tray showing the time as 12:44 pm.

RT ticketing system: Queue status

The screenshot shows a Mozilla Firefox browser window displaying the RT ticketing system interface for 'cslab.ntua.gr'. The page title is 'Found 24 tickets'. The user is logged in as 'gef'. The interface includes a search bar, a navigation menu, and a table of tickets.

#	Subject Requestor(s)	Status Created	Queue Last Contact	Owner Last Updated	Priority Left
10949	Guide for new HW sites HG-04-CTI-CEID@hellasgrid.gr, hg-02-iasa@hellasgrid.gr, hg-03-auth@hellasgrid.gr, hg-05-forth@hellasgrid.gr, kkoum@gmet.gr	open 5 months	HG-01-GRNET 2 weeks	azisi 72 min	0
12297	GGUS-Ticket-ID: #6371 assigned to Generic Deployment - LB Deployment in LCG... support@ggus.org	new 6 days	HG-01-GRNET -	Nobody 5 hours	0
12375	Re-Configuration kkoum@gmet.gr	new 6 hours	HG-01-GRNET -	Nobody 6 hours	0
12368	\$ATLAS_SW_DIR problems cpapachr@ics.forth.gr	open 11 hours	HG-01-GRNET 8 hours	vkoukis 8 hours	0
12209	Storage Element upgrade se.dpm/dCache kkoum@gmet.gr	open 2 weeks	HG-01-GRNET -	kyrginis 20 hours	0
12350	upgrade to 2.7.0 kkoum@gmet.gr	open 3 days	HG-01-GRNET 3 days	Nobody 3 days	0
12339	GlueCEPolicyMaxRunningJobs: 0 kkoum@gmet.gr	new 4 days	HG-01-GRNET -	kyrginis 4 days	0
12340	dCache vs. DPM kkoum@gmet.gr	new 4 days	HG-01-GRNET -	Nobody 4 days	0
12334	[Fwd: [HellasGrid User-request] Νέα αίτηση] kkoum@gmet.gr	new 4 days	HG-01-GRNET -	Nobody 4 days	0
810	Να βρεθεί υποστήριξη και upgrades για airco/ups κλπ. fotis@mail.cern.ch, kkoum@gmet.gr	open 1 years	HG-01-GRNET 4 days	kkoum 4 days	0
12319	[Fwd: Re: Cisco 3750 Stack - error messages] azisi@cslab.ece.ntua.gr	new 5 days	HG-01-GRNET -	Nobody 5 days	0

Find: vangelis Find Next Find Previous Highlight Match case

https://rt.cslab.ece.ntua.gr/Ticket/Display.html?id=12209

RT system: managing a ticket

The screenshot shows a Mozilla Firefox browser window with the following details:

- Address Bar:** `https://rt.cslab.ece.ntua.gr/Ticket/Display.html?id=12368`
- Page Title:** #12368: \$ATLAS_SW_DIR problems
- Browser Tabs:** LCG Generic Installation & Configuration Gu..., RT at a glance, Found 24 tickets, #12368: \$ATLAS_SW_DIR problems
- Page Content:**
 - History:** 12368: \$ATLAS_SW_DIR problems (open)
 - Display mode:** [Brief headers] [Full headers]
 - Message #1:** Tue Feb 07 12:40:02 2006 **gernot.krobath@Physik.Uni-Muenchen.DE** - Ticket created [Reply] [Comment] Download (untitled) 480b
 - Subject:** \$ATLAS_SW_DIR problems
 - Date:** Tue, 07 Feb 2006 11:37:21 +0100
 - To:** hg-01-grnet@hellasgrid.gr
 - From:** Gernot Krobath <gernot.krobath@Physik.Uni-Muenchen.DE>

Hi,

I ran a check-site at ce01.ariagni.hellasgrid.gr and I got the following errors:

```
du: /opt/exp_soft/atlas/*: No such file or directory
du: /opt/exp_soft/atlas/software/*: No such file or directory
ls: /opt/exp_soft/atlas/software: No such file or directory
ls: /opt/exp_soft/atlas/software/*/setup.sh: No such file or directory
```

probably your \$ATLAS_SW_DIR is not set correctly.
Please fix this and tell me as soon as it's been fixed.

Sincerely yours,
Gernot Krobath.
 - Message #2:** Tue Feb 07 13:24:15 2006 **vkoukis** - Correspondence added [Reply] [Comment] Download (untitled) 783b
 - RT-Send-CC:** HG-05-FORTH@hellasgrid.gr
 - Quoted text:**

```
> [gernot.krobath@Physik.Uni-Muenchen.DE - Tue Feb 07 12:40:02 2006]:
>
> Hi,
>
> I ran a check-site at ce01.ariagni.hellasgrid.gr and I got the following
> .....
```
- Search Bar:** Find: vangelis [Find Next] [Find Previous] [Highlight] [Match case]
- Taskbar:** start, 2 Firefox, Inbox for Fotis@mail..., Microsoft PowerPoint ...

Notorious Tickets of the past

- #62: Evaluating GPFS suitability for our case
- #304: Optimizing configuration of CISCO switch stack
- #809: Document startup, boot, halt shutdown procedures
- #888: Daily check-list of the hg-01-grnet node
- #942: GPFS debugging & tiobench tests
- #2078: mprime and kernel-compiles, cpu temperature tests
- #4804: Heavy stress tests of HG-01-GRNET (minor repairs)
- #7200: Getting mpirun to work correctly with torque
- #7293: VPN architecture and OOB management
- #8460: Unscheduled downtime due to power
- #8617: Unscheduled downtime due to air-condition failure

- Evaluating GPFS suitability for our case
- Optimization configuration of CISCO switch stack
- Document startup, boot, halt and shutdown procedures
- Build a daily check-list for the HG-01-GRNET node
- GPFS debugging & tiobench tests
- kernel-compiles and cpu temperature tests with mprime
- Heavy stress tests of HG-01-GRNET (with minor repairs)
- Unscheduled downtime due to power
- Unscheduled downtime due to air-condition failure

Get ready to coordinate with the plumber... ☺

Putting it all together

23 IBM x335 xSeries Servers



ISCSI Cisco SN 5428

8 IBM x335 xSeries
Servers SAN Attached

IBM 3534
FC Switches

TSM
Backup
Server

Management
Server


FastT900



3583 – L72 Library



HellasGrid I, Isabella



- High-performance parallel, scalable file system for Linux/AIX cluster environments
- Full Load Balancing: Shared-disk filesystem where every cluster node can have concurrent read/write access to (**meta-**)data
- Fail-over: High availability through automatic recovery from node and disk failures

Monitoring of HG-01-GRNET

- hg-01-grnet@hellasgrid.gr, [hwinfo\(...\)@rt.cslab.ntua.gr](mailto:hwinfo(...)@rt.cslab.ntua.gr)
- Tivoli Storage Manager – TSM backup interface
- CISCO switch stack
- CISCO iSCSI switches
- Cluster Systems Management (CSM) utilities, RSA, RCM
- GPFS and filesystems monitors
- Storage Area Network GUI client

Another “minor” cluster waiting 4 u



Q & A

