

Existing EGEE Metrics

*Piotr Nyczyk
CERN IT/GD*

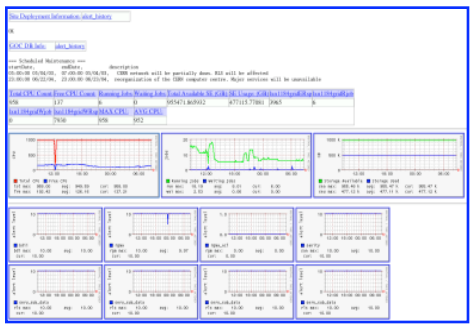
*Joint OSG and EGEE Operations Workshop - 3
Abingdon, 27 - 29 September 2005*

- **Motivation**
- **Integration of monitoring information**
- **Metrics calculation**
- **Graphs**
- **Links and feedback**

- **We need to measure the effectiveness of operations:**
 - CIC-on-duty impact
 - ROCs effectiveness to fix problems
 - Sites responsiveness
- **Various monitoring tools show many different views/ aspects - but no conclusions:**
 - “Is it any better now than it was one week ago?”
 - “Is the improvement fast or slow?”
- **Consequence: we need a simple number for a site, region, whole grid that measures “how good is the grid”**
- **BTW. Users (experiments) don’t care about our internal procedures - they look at the resources which they can use (or cannot) - consequence: scheduled downtime is unavailability**

Site Name	Status	Site Name	Status	Site Name	Status
BRISTOL-LCG2	OK	CE-FRAN-LCG2	OK	CSN1-LCG2	OK
EPFL-LCG2	OK	DFK-FRAN-LCG2	OK	EPFL-LCG2	OK
INFN-BO-LCG2	OK	INFN-BO-LCG2	OK	INFN-BO-LCG2	OK
INFN-BO-LCG2	OK	INFN-BO-LCG2	OK	INFN-BO-LCG2	OK
INFN-BO-LCG2	OK	INFN-BO-LCG2	OK	INFN-BO-LCG2	OK

1. GIS Monitor



2. GIS Monitor graphs

Site	Test Name	Test Date	Test Status	Test Duration	Test Results
pp.frac.uk	Site Functional Tests	2005-01-14	Pass	10:00:00	OK
pp.frac.uk	Site Functional Tests	2005-01-14	Pass	10:00:00	OK
pp.frac.uk	Site Functional Tests	2005-01-14	Pass	10:00:00	OK

3. Sites Functional Tests

Site Information - RAL-LGG2

Site Name: RAL-LGG2

Site Address: RAL, Harwell, Didcot, Oxfordshire, UK

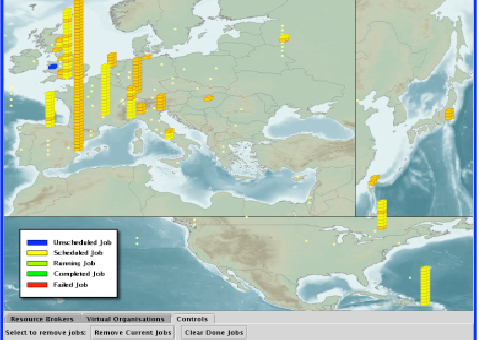
Site Contact: RAL-LGG2

Site Description: RAL-LGG2

4. GOC Data Base

Site	Description	Start Date	End Date
WIRTH-LCG2	Upgrade of hardware and LCG software upgrade	2005-01-14 13:45:00	2005-01-21 23:59:00
UCL-REP	Replacement of CE and upgrade of job & LCG 2.2.0	2005-01-12 18:00:00	2005-01-17 23:59:00
INSP-LAL	Upgrade to LCG 2.2	2005-01-15 08:00:00	2005-01-17 08:00:00

5. Scheduled Downtimes



6. Live Job Monitor

Job Monitoring view History

VO Name	Jobs Running	Jobs Queued	Jobs Total
pp.frac.uk	11	0	11
pp.frac.uk	2	0	2
pp.frac.uk	13	0	13
pp.frac.uk	10	0	10
pp.frac.uk	25	0	25
TOTAL	61	0	61

7. Gridce - VO view

GridCE - fabric view

Site: pp.frac.uk

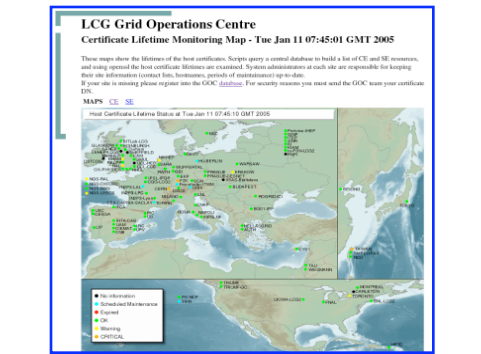
Site Name: pp.frac.uk

Site Address: pp.frac.uk

Site Contact: pp.frac.uk


Site Description: pp.frac.uk

8. Gridce - fabric view



9. Certificate Lifetime Monitor

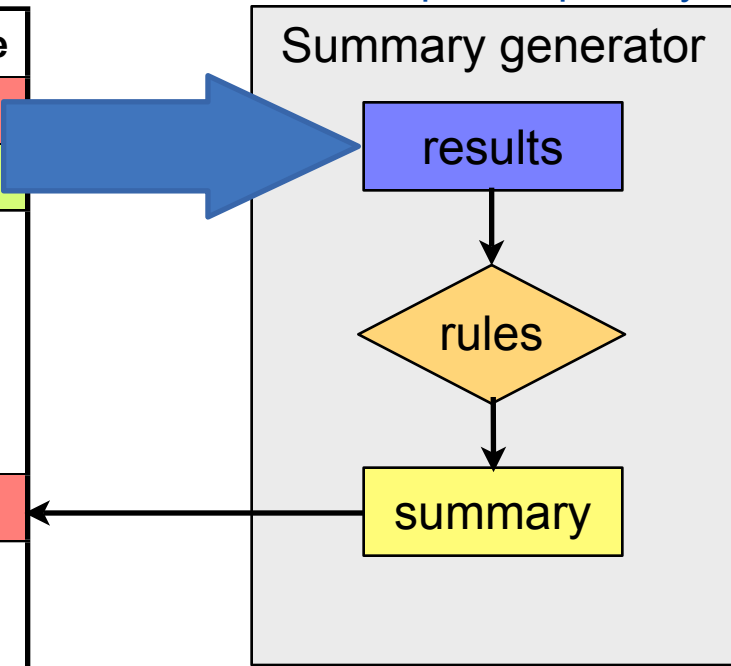
- Monitoring information from various tools is collected in R-GMA archiver
- Summary generator calculates overall status of each monitored object (site, CE, ...) - update: 1h
- Metric generator calculates numerical value for each monitored object + aggregation (CE → site → region → grid) - update: 1 day



- Summary generator:**

- performs “reduce” operation on results for all VOs, tests, sites
- uses rules defined by choice of critical tests (FCR) and by relations between objects (site owns 1 or more CEs, etc.)
- summaries are also results → a hierarchy of rules, example:
 - SFT critical tests → sft-summary
 - gstat tests → gstat-summary
 - sft-summary + gstat-summary → site-summary
- operation is repeated every 1 hour: hourly snapshot for all sites, 24 snapshots per day

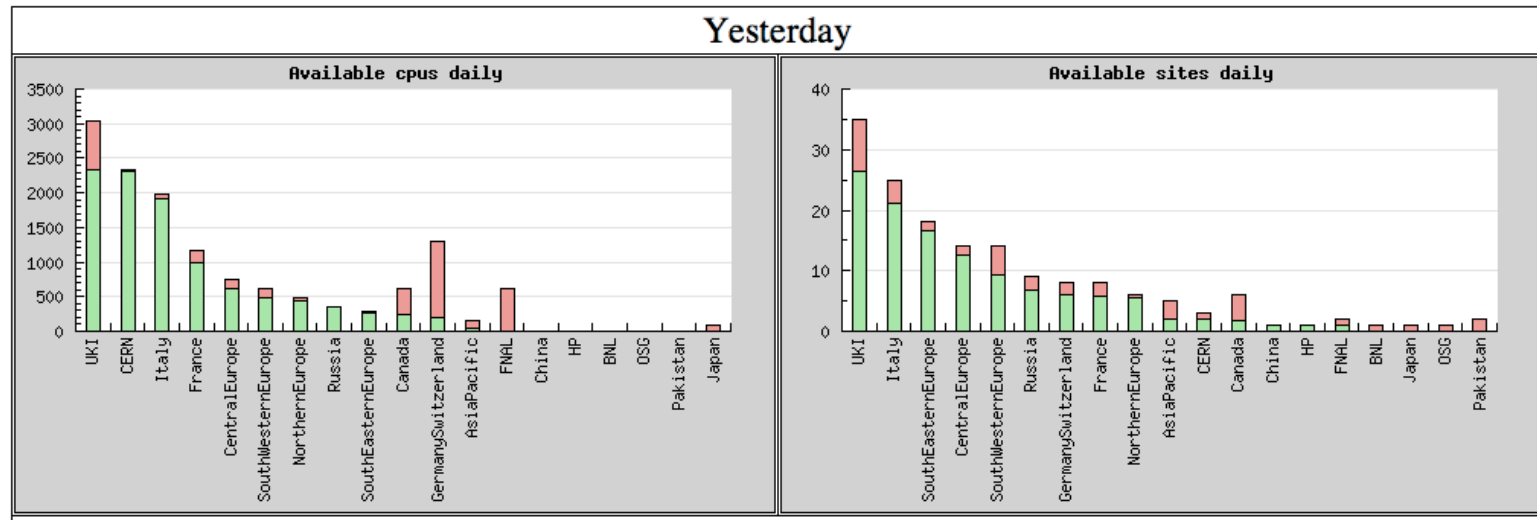
VO	testName	nodeName	status	time
dteam	sft-job	ce1.badsite.org	ERROR	T1
-	gstat-GIIS	BADSITE-LCG2	OK	T2
.
.
.
dteam	site-summary	BADSITE-LCG2	ERROR	T3
.
.
.



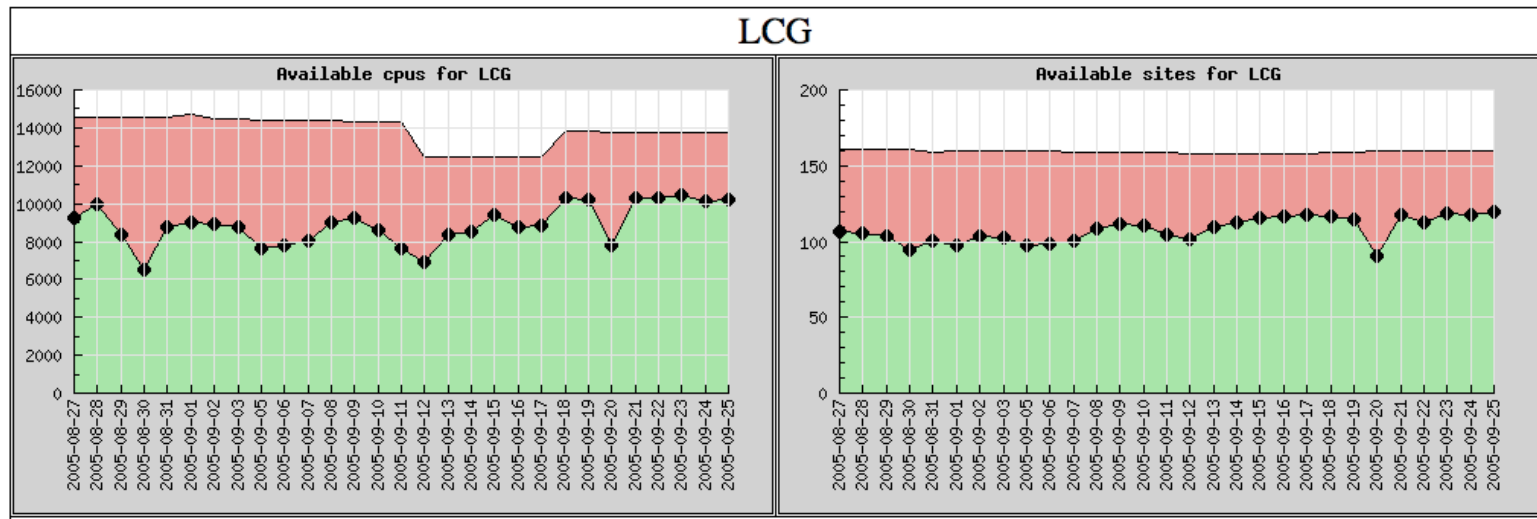
- **Sites availability:**
 - for each site: 24 summaries per day, each has two possible values 0 - ERROR or 1 - OK
 - summary values integrated (averaged) for the whole day ⇒
 - 0 - site was failing tests for the whole day
 - 0.75 - site was passing tests for 18 hours and failing for 6 hours
 - 1 - site was OK for the whole day
 - metric values for sites in a single region are added together to calculate the values for the whole region
 - metric values of all regions are summed to get overall metric for the grid
- **CPUs availability**
 - similar process, **but:** site value is multiplied by the number of CPUs in that site
 - example: site A has 1000 CPUs but the daily site availability metric was 0.80 ⇒ CPUs availability metric is 800
 - interpretation: How many cpus were **effectively** provided by a site/ region/grid during the day?

- **Long term metrics:**
 - metric results are archived in R-GMA table (MetricData)
 - weekly, monthly, *quarterly* metric is calculated by averaging daily metric for site/region/grid
- **Potential CPUs availability metric:**
 - indicates how many CPUs could be provided by the site/region/grid if all tests were always passed
 - is calculated by taking 1 as the value of site availability metric for all sites

Averaged metrics



History metrics (monthly)



- **Prototype metrics report is available:**
<https://lcg-sft.cern.ch:9443/sft/metrics.html>
- **After announcing the prototype report I received useful feedback and suggested improvements:**
 - tables with the actual value series
 - graphs/tables for individual sites - useful for ROCs to see how sites perform
- **Further discussion and comments are welcome...**