

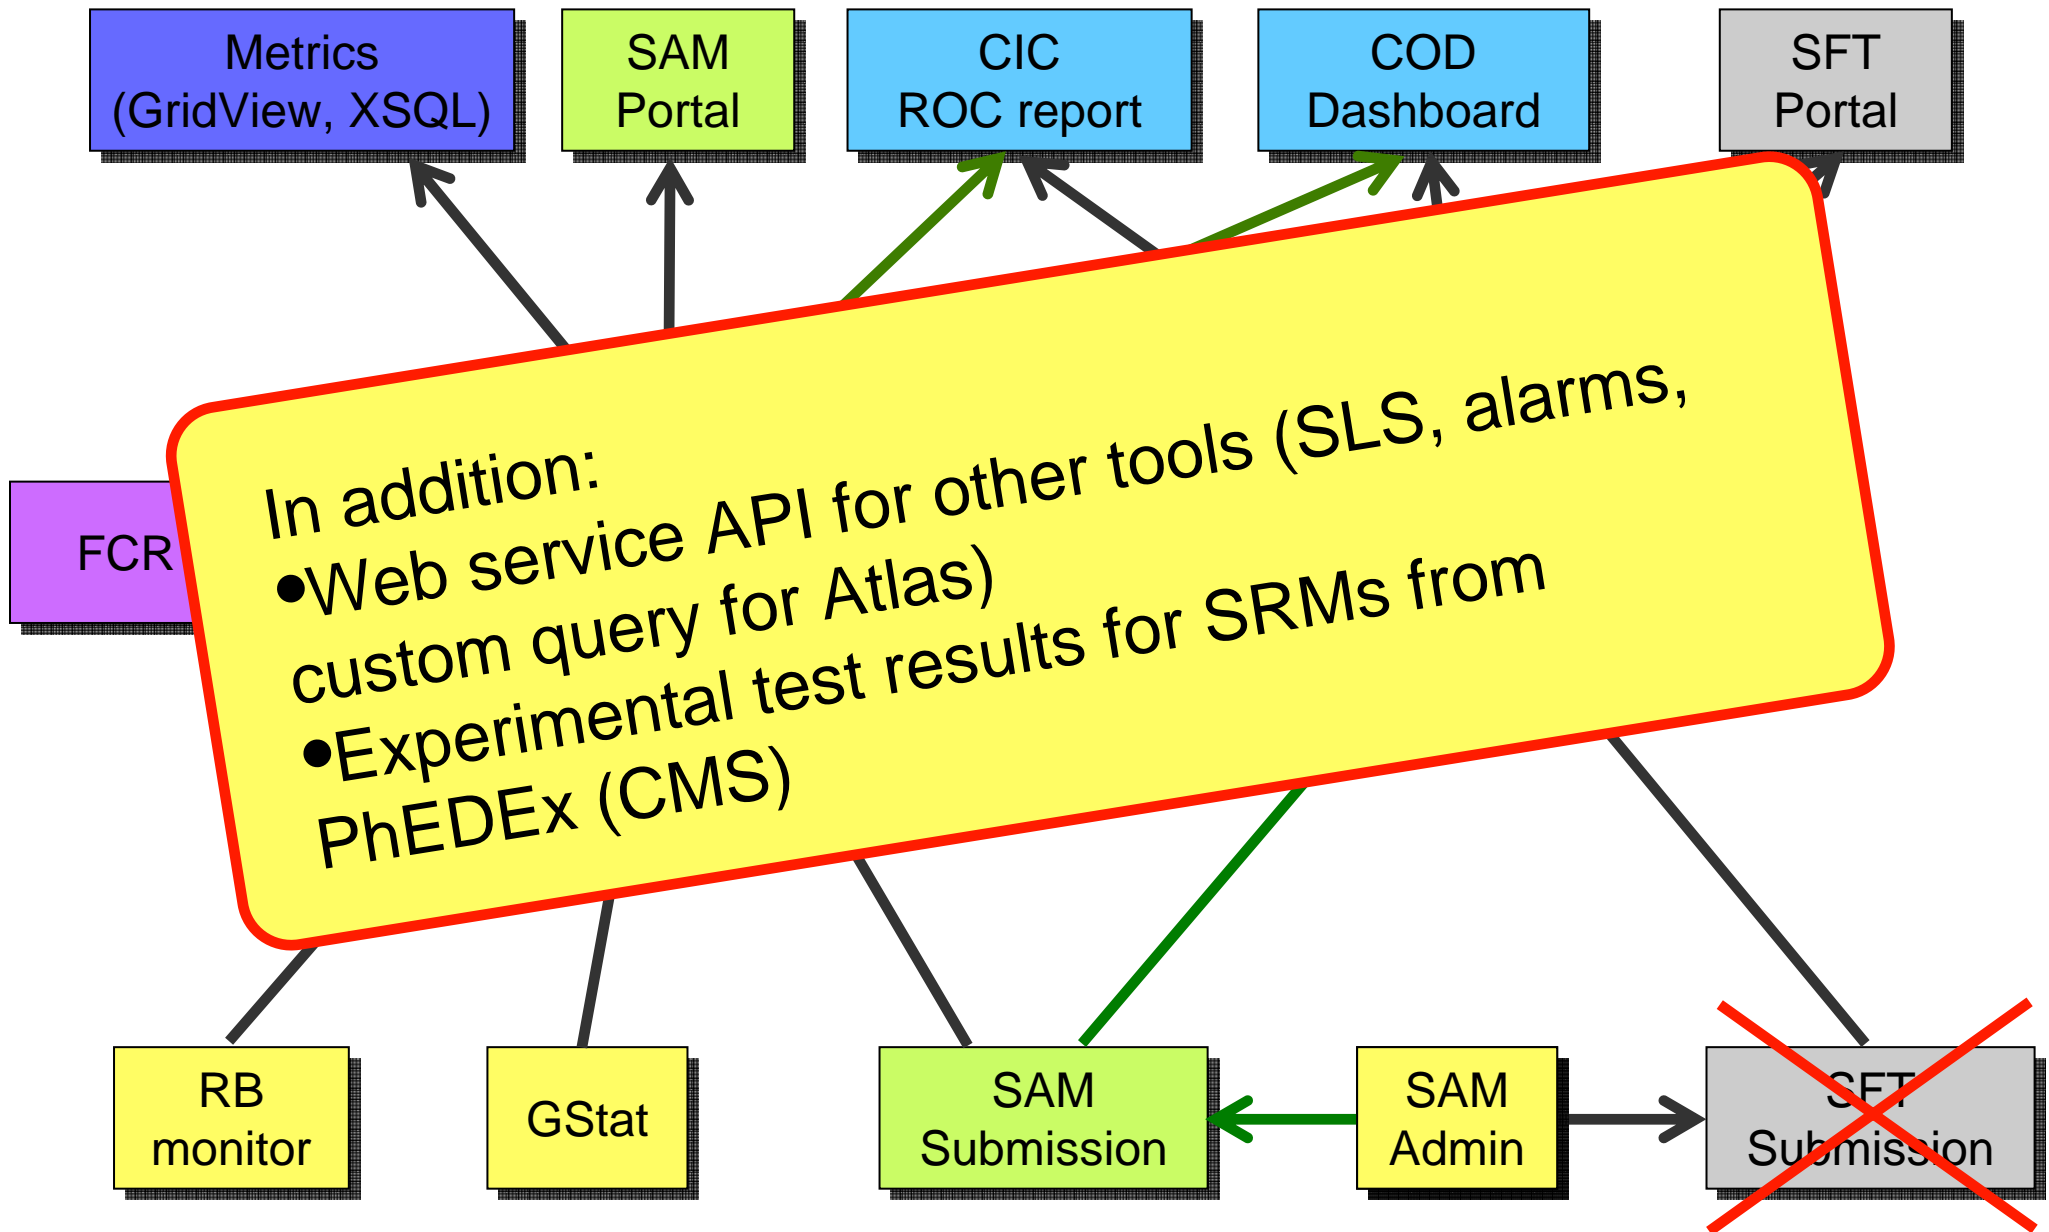
# SAM Status



Piotr Nyczyk, CERN IT/GD

Grid Deployment Board  
08 November 2006, Geneva

- Deployment status
- Sensors status
- Existing tests
- VO specific tests
- Jobwrapper tests
- Availability metrics
- Open issues



Service	Responsible	Status
SRM	P. Nyczyk	Done
LFC	J. Casey	Done
FTS	P.Nyczyk	Done (static list of files for transfer tests)
CE	P. Nyczyk	Done
RB	D. Kant	Done
RB-pasv	S. Andreozzi	No news
gLite WMS	?	Certification tests migrated to SAM in CTB
Top-level BDII	M. Tsai	Done
Site BDII	M. Tsai	Done
MyProxy	M. Litmaath	To be done (simple tests)
VOMS	V. Venturi	Sensor written, to be integrated
R-GMA	?	Script written, sensor has to be created
Tier-1 DB	?	Suggested by Dirk Duellmann, to be done

- job submission - UI->RB->CE->WN chain
- version of CA certificates installed (on WN!)
- version of software middleware (on WN!)
- broker info - checking *edg-brokerinfo* command
- UNIX shells environment consistency (BASH vs. CSH)
- replica management tests - using *lcg-utils*, *default SE* defined on WN and a selected “central” SE (3-rd party replication)
- accessibility of experiments software directory - environment variable, directory existence
- accessibility of VO tag management tools
- other tests: R-GMA client check, Apel accounting records

- storing file from the UI - using *lcg-cr* command with LFC registration
- getting file back to the UI - using *lcg-cp* command
- removing file - using *lcg-del* command with LFC de-registration

- LFC
  - directory listing - using *lfc-ls* command on /grid
  - creating file entry in /grid/<VO> area
- FTS
  - checking if FTS is published correctly in the BDII
  - channel listing - using *glite-transfer-channel-list* command with ChannelManagement service
  - transfer test (in development):
    - submitting transfer jobs between SRMs in all Tier0 and Tier1 sites (N-N testing)
    - checking the status of jobs
    - Note! The test is relying on availability of SRMs in sites

- GStat:
  - site-BDIs: accessibility (response time), sanity checks (partial Glue schema validation)
  - top-level BDIs: accessibility (response time), reliability of data (number of entries)
- RB:
  - jobs submitted through all important RBs to selected “reliable” CEs
  - measuring time of matchmaking



- LHCb
  - already submits jobs, but using old SFT framework (only CEs, need to migrate to SAM)
  - VO specific tests (Dirac installation)
- Atlas
  - jobs for all sensors submitted from SAM UI with Atlas credentials
  - standard set of tests
- CMS
  - account on SAM UI created, sample jobs sent, no regular submission yet
  - VO is using OPS results in FCR
- Alice - no information

- Requested by experiments - motivation:
  - SAM/SFT jobs are not reaching all WNs, thus not detecting broken WNs which can lead to unacceptable failure rate
- Simplified set of tests executed on WNs by the CE wrapper with each grid job
- Will potentially reach all WNs in relatively short time
- Test results will be passed to the job but also published to the central DB
- Core scripts have to be installed on CE and WNs (will become part of the release)
- Set of tests will be updated centrally (signed tarball installed in software area)


- Gains for operations:
  - infrastructure description with unique identification of WNs:
    - relation between CEs, batch queues, WNs
    - detection of monitoring queues pointing to “carefully” selected WN
    - counting WNs at the site without risk of double counting due to shared batch farms
  - detection of sites with broken WNs: basic fabric monitoring for small sites
- Status:
  - scripts, tests, transport layer ready and certified on CTB
  - deployment on PPS should start soon
  - visualisation and data processing tools have to be written (simple view at CERN, data analysis by team from Wuppertal?)

- gLite WMS
  - tests from Certification Test Suite were migrated to SAM in Certification TestBed (but mostly stress tests)
- MyProxy sensor - ???
- VOMS sensor
  - some work done by Valerio Venturi, needs to be integrated
  - additional tests are being developed by CTB team
- R-GMA registry
  - sensor has to be written, but not very urgent
- Tier1 DB
  - need for monitoring of selected Oracle DB services at Tier1 centres
  - sensor has to be written almost from scratch
- R-GMA

Status of node N =  $\bigwedge_{t \in \text{CriticalTests}} \text{TestResult}(N,t)$

$\wedge$  = boolean AND  
 $\vee$  = boolean OR


Status of central service C =  $\bigvee_{N \in \text{instances}(C)} \text{Status}(N)$



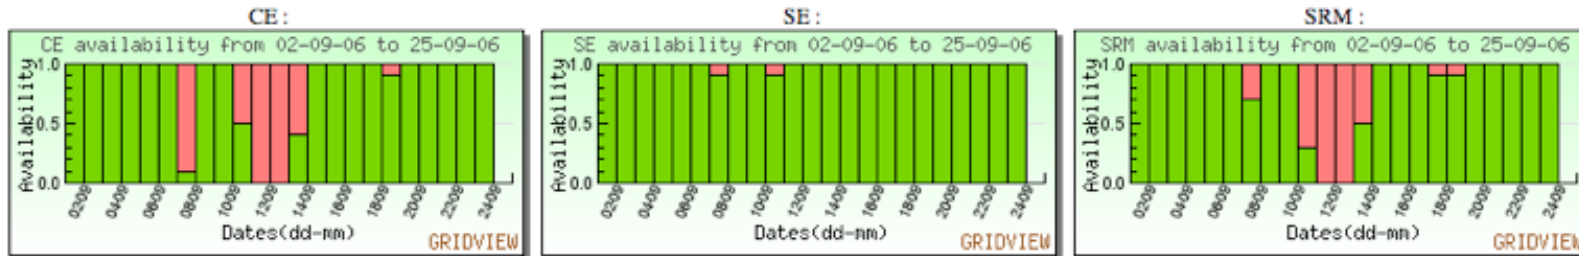
Everything is calculated for each VO that defined critical tests in FCR  
 Results make sense only if VO submits tests!!!

- Service and site status values are recorder every hour (24 snapshots per day)
- Daily, weekly, monthly availability is calculated using integration (averaging) over the given period
- Scheduled downtime information from GOC DB is also recorder and integrated for comparisons
- Details of the algorithm on GOC:  
[http://goc.grid.sinica.edu.tw/gocwiki/SAM\\_Metrics\\_calculation](http://goc.grid.sinica.edu.tw/gocwiki/SAM_Metrics_calculation)

## Overall Service Availability for site CERNPR : Daily Report



## Individual Service Availability for site CERNPR : Daily Report




## Service Instance Availability for site CERNPR : Daily Report



## Tier-1/0 Site Availability : Daily Report

(Click on the Graph below to see Availability of Individual Services at the Site)



- NOTE. LOW SCORE OF BNL AND IN2PCC CAUSED BY SRM failures



Availability of T0 and T1 sites including downtime and unscheduled downtime.

http://lcg-sam.cern.ch:8080/sqlldb/site\_avail.xsql

Workbook1

Availability of sites for ' OPS ' VO including downtimes

Date	BNL-LCG2	CERN-PROD	FZK-LCG2	IN2P3-CC	INFN-T1	RAL-LCG2	SARA-MATRIX	TRIUMF-LCG2	Taiwan-LCG2	pic																					
1/9/06	0	0	1	1	0	0	0.9	0	0.1	0	0	1	1	0	0	0.9	0	0.1	1	0	0	0.5	0	0.5	0	0.5					
2/9/06	0	0	1	0.9	0	0.1	0	0	1	0	0	1	1	0	0	0.7	0	0.3	1	0	0	0.8	0	0.2	0	0	1	1	0	0	
3/9/06	0	0	1	1	0	0	0	0	1	0	0	1	1	0	0	0.3	0	0.8	1	0	0	0.9	0	0.1	0	0	1	1	0	0	
4/9/06	0	0	1	1	0	0	0.3	0	0.8	0	0	1	1	0	0	0.6	0	0.4	0.88	0	0.13	0.4	0	0.6	0	0	1	0.8	0	0.2	
5/9/06	0	0	1	1	0	0	0.3	0	0.7	0	0.1	0.9	1	0	0	0.7	0	0.3	0.83	0	0.17	0.8	0	0.2	0	0	1	0.7	0	0.3	
6/9/06	0	0	1	1	0	0	0.5	0	0.5	0	1	0	0.6	0	0.4	0.5	0	0.5	0.5	0	0.5	0.3	0	0.8	0	0	1	1	0	0	
7/9/06	0	0	1	1	0	0	0.9	0	0.1	0	1	0	0	0	0.5	0	0.5	0.71	0	0.29	0.7	0	0.3	0	0	1	0.4	0	0.6		
8/9/06	0	0	1	1	0	0	1	0	0	0	0.8	0.2	0.9	0	0.1	0.8	0	0.3	0.79	0	0.21	0.3	0	0.7	0	0	1	1	0	0	
9/9/06	0	0	1	0.3	0	0.7	0.5	0	0.5	0	0.5	0.5	0.8	0	0.2	0.2	0	0.8	0.29	0	0.71	0.2	0	0.8	0	0	1	0.7	0	0.3	
10/9/06	0	0	1	1	0	0	0.6	0	0.4	0	0	1	1	0	0	0.7	0	0.3	1	0	0	0.9	0	0.1	0	0	1	1	0	0	
11/9/06	0	0	1	0.9	0	0.1	0	0	1	0	0	1	1	0	0	0.8	0	0.3	1	0	0	0.3	0	0.7	0	0	1	1	0	0	
12/9/06	0	0	1	0.3	0	0.7	0.3	0	0.7	0	0	1	1	0	0	0.3	0	0.7	1	0	0	0.5	0	0.5	0	0	1	0.7	0	0.3	
13/9/06	0	0	1	0	0	1	0.3	0	0.7	0	0	1	1	0	0	0.5	0	0.5	0.92	0	0.08	0.7	0	0.3	0	0	1	1	0	0	
14/9/06	0	0	1	0.2	0	0.8	0.8	0	0.3	0	0.5	0.5	1	0.5	0	1	0	0	1	0	0	0.8	0	0.3	0	0	1	1	0	0	
15/9/06	0	0	1	1	0	0	0.8	0	0.3	0	0	1	1	0.8	0	1	0	0	1	0	0	0.6	0	0.4	0	0	1	1	0	0	
16/9/06	0	0	1	1	0	0	0.9	0	0.1	0	0	1	1	0	0	1	0	0	1	0	0	0.2	0	0.8	0	0	1	1	0	0	
17/9/06	0	0	1	1	0	0	1	0	0	0	0	0	1	1	0	0	0	1	0	0	0	0.7	0	0.3	0	0	1	1	0	0	
18/9/06	0	0	1	1	0	0	1	0	0	0	0	0	1	1	0	0	0	1	0	0	0	0.6	0	0.4	0	0	1	1	0	0	
19/9/06	0	0	1	0.9	0	0.1	0.9	0	0.1	0	0	1	1	0	0	1	0	0	0.92	0	0.08	0.2	0.6	0.3	0	0	1	0.6	0	0.4	
20/9/06	0	0	1	0.9	0	0.1	0.3	0	0.7	0	0	1	1	0	0	0.9	0	0.1	0.58	0	0.42	0	0.7	0.3	0	0	1	0.7	0	0.3	
21/9/06	0	0	1	1	0	0	0.3	0	0.7	0	0	1	0.8	0	0.3	0.1	0	0.9	0.96	0	0.04	0	0	1	0	0	1	1	0	0	
22/9/06	0	0	1	1	0	0	1	0	0	0	0	1	1	0	0	0.8	0	0.3	1	0	0	0	0	1	0	0	1	1	0	0	
23/9/06	0	0	1	1	0	0	0.6	0	0.4	0	0	1	0.6	0	0.4	0.6	0	0.4	0.96	0	0.04	0	0	1	0	0	1	1	0	0	
24/9/06	0	0	1	1	0	0	0.9	0	0.1	0	0	1	0	0	1	0.4	0	0.6	1	0	0	0	0	1	0	0	0.3	0.8	0.9	0	0.1
Average	0		0.85			0.58			0		0.9		0.67		0.89		0.42														0.87

For each site, there are [0-1] values in 3 columns representing:

- GREEN: the availability of the site.
- ORANGE: the scheduled downtime.
- RED: the unscheduled downtime.

Sheet1

- All sensors have to be reviewed and fixed:
  - check if tests reflect real usage (experiments)
  - avoid dependencies on central services and third party services if possible
  - increase reliability of results (resistant to any other failures not related to site configuration)
  - increase tests verbosity (make easier to find real problem - site debugging))
- Missing sensor/tests have to be written
- All tests should be well documented (TestDef inline doc + Wiki)
- Jobwrapper tests should be put in production (simple display or data export needed)
- Availability metric should be reviewed and fixed - *is the current one the way to go?*
- Metric calculation for aggregates of sites