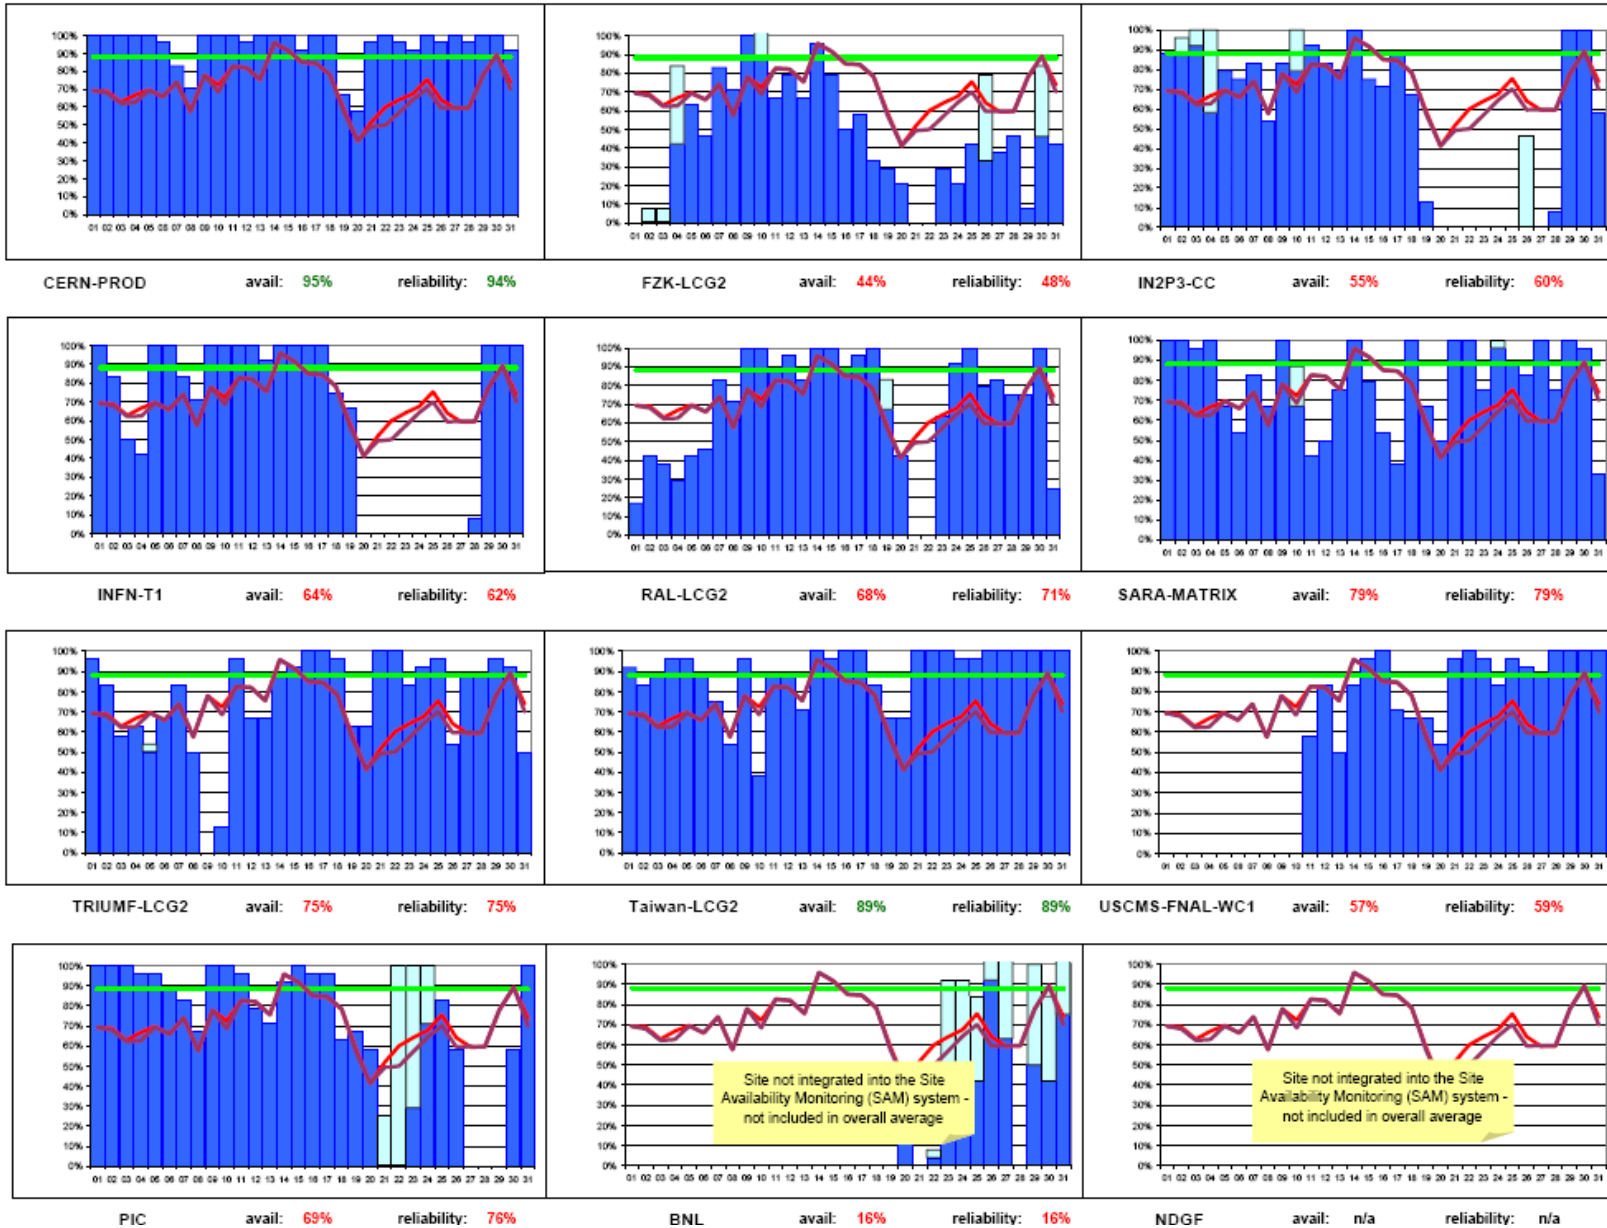


October site availability reports from Tier-1 centres










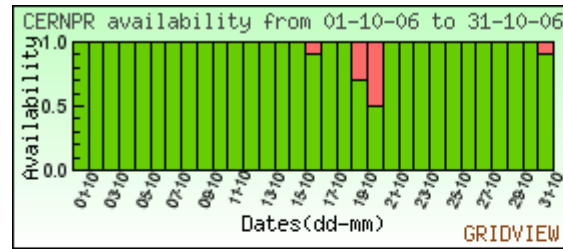
MB 7-Nov at CERN

G. Merino


October availability report




October SAM results from GridView




Larger differences Availability report vs. GridView




IN2P3-CC avail: 55% reliability: 60%




IN2P3




INFN-T1 avail: 64% reliability: 62%




CNAF




TRIUMF-LCG2 avail: 75% reliability: 75%



TRIUMF



Taiwan-LCG2 avail: 89% reliability: 89%



ASGC


IN2P3 report

Period 19-28 Oct:


- Almost 0% availability according to the report
- GridView and Fabio's check of detailed SAM tests shows a different reality

Date	Virtual Organization			SRM: ccsr.in2p3.fr			CE:cclgceli02.in2p3.fr			CE:cclgceli03.in2p3.fr		
	dteam	ops	best	Successes	Failures	Daily Metric	Successes	Failures	Daily Metric	Successes	Failures	Daily Metric
19/10/2006	0,38	0,13	0,38	14	26	35%	0	7	0%	9	8	53%
20/10/2006	0,88	0,00	0,88	23	16	59%	11	8	58%	16	8	67%
21/10/2006	1,00	0,00	1,00	4	0	100%	3	1	75%	3	0	100%
22/10/2006	1,00	0,00	1,00	2	0	100%	0	1	0%	1	0	100%
23/10/2006	1,00	0,00	1,00	20	4	83%	15	0	100%	13	0	100%
24/10/2006	0,88	0,00	0,88	19	5	79%	16	2	89%	14	1	93%
25/10/2006	1,00	0,00	1,00	10	5	67%	11	0	100%	9	1	90%
26/10/2006	1,00	0,00	1,00	20	4	83%	14	1	93%	15	0	100%
27/10/2006	0,88	0,00	0,88	21	3	88%	10	2	83%	10	2	83%
28/10/2006	1,00	0,08	1,00	20	0	100%	9	1	90%	8	1	89%
29/10/2006	1,00	1,00	1,00	24	0	100%	12	0	100%	12	0	100%
30/10/2006	1,00	1,00	1,00	17	0	100%	11	0	100%	12	0	100%
31/10/2006	0,83	0,58	0,83	20	0	100%	4	5	44%	3	5	38%

Overall Service Availability for site IN2PCC : Daily Report



Individual Service Availability for site IN2PCC : Daily Report



Operational issues: dCache

- Several sites report about the “hanging gridftp doors” dCache scalability problems causing unavailability.
- Some of them report about corrective actions taken:
 - **NIKHEF**
 - Improvement of problem-detection scripts
 - Upgrade to dCache v1.7 → situation improved
 - **RAL**
 - reduced tcp window size on GridFTP doors → situation improved
 - **FZK**
 - Solve inconsistency: srm from gLite was “too” new
 - Tune several dCache parameters (max logins, number of streams per client) → situation slightly improved
 - Automatically restarting the gridftp doors → situation improved for experiments, but not for SAM tests
 - dCache upgrade from v1.6.6.0 to v1.6.6.5 → situation did not improve

Operational issues: CE

- **Stability issues**

- **RAL:** upgraded CE to a 2xCPU, 2xCore with 4GB RAM → improved load-related problems
- **PIC:** suffering a lot from users submitting through many RBs and killing the CE. Planning to deploy a 2nd CE (redundant service) on a more powerful hardware.

- **Torque security vulnerability (20-Oct, Friday)**

- **RAL:**
 - Patch applied in less than 1 day, but CE SAM tests still failing for about 3-4 days later. Reason not understood.
 - After this patching, problems appeared in the CE (a WN can stop the server scheduling jobs). Still investigating.
- **FZK:** PBSPro released the patch on Monday. Queues closed the whole weekend.
- **PIC:** Patch applied in “urgent” mode, created other problems in the WNs configuration that caused CE intermittent unavailability for some days.


SAM issues: site aggregation

- SARA-NIKHEF is a Tier-1 centre formed by two LCG sites
 - The CE service from SARA-NIKHEF Tier-1 is 90% at NIKHEF, so SARA CE unavailability is not representative of the service
- **ISSUE:** SAM not able to manage Tier-1s made of an “aggregation” of sites


SAM issues: errors summary info

- Suggestion from IN2P3: It would very useful if SAM reports were easier to interpret and allowed site managers to easily spot:
 - **a)** What service (CE, SRM, BDII) is making the site unavailable
 - *...this is already addressed by the gridview tool:*

Overall Service Availability for site IN2PCC : Daily Report



Individual Service Availability for site IN2PCC : Daily Report



- **b)** What tests (get, put, JS, JL, etc.) the individual service is frequently failing
 - *... may be producing a table with the “number of failures for each test-type” in the selected time period?*

SAM issues: OPS vs DTEAM

- In the October availability report, only TRIUMF and Taiwan have some “DTEAM availability” at the beginning of the month.
- CNAF
 - Reports that had issues with configuring the OPS VO in the SE service. DTEAM tests ok.
 - As of yesterday, SRM service still does not show up in the SAM OPS VO page for CNAF.

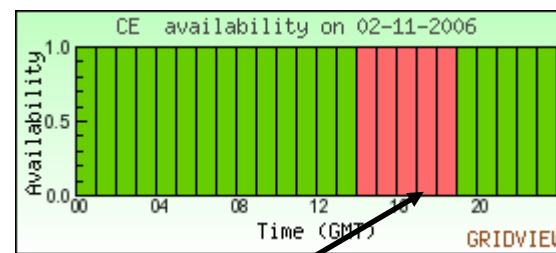
SAM issues: false positives

Timeout when executing test CE-sft-lcg-rm after 600 seconds!

Contact: SAM Development Mailing List ([<same-devel@cern.ch>](mailto:same-devel@cern.ch))

Last modified: Nov 02 2006, 09:02

Test	Result	Details
CE-sft-lcg-rm-gfal	OK	SAM SFT
CE-sft-lcg-rm-cr	OK	SAM SFT
CE-sft-lcg-rm-cp	OK	SAM SFT



				3.0.2	ok	ok	ok	ok
10	02-Nov-2006 18:57:59	OK	ok	3.0.2	ok	ok	ok	ok
11	02-Nov-2006 13:56:50	ERROR	ok	3.0.2	ok	ok	ok	error
12	02-Nov-2006 11:51:29	OK	ok	3.0.2	ok	ok	ok	ok
13	02-Nov-2006 10:53:43	ERROR	error	3.0.2	ok	ok	ok	ok
14	02-Nov-2006 08:49:05	OK	ok	3.0.2	ok	ok	ok	ok
15	02-Nov-2006 07:49:48	OK	ok	3.0.2	ok	ok	ok	ok
16	02-Nov-2006 05:49:51	OK	ok	3.0.2	ok	ok	ok	ok
17	02-Nov-2006 03:49:04	OK	ok	3.0.2	ok	ok	ok	ok
18	02-Nov-2006 02:49:47	OK	ok	3.0.2	ok	ok	ok	ok
19	02-Nov-2006 00:51:58	ERROR	error	3.0.2	ok	ok	ok	ok

- icg-rm test failing in replication to CERN step. Looks like false positive.
- Next SAM test takes 5 hours to run.

=> Unrealistic 21% site unavailability that day

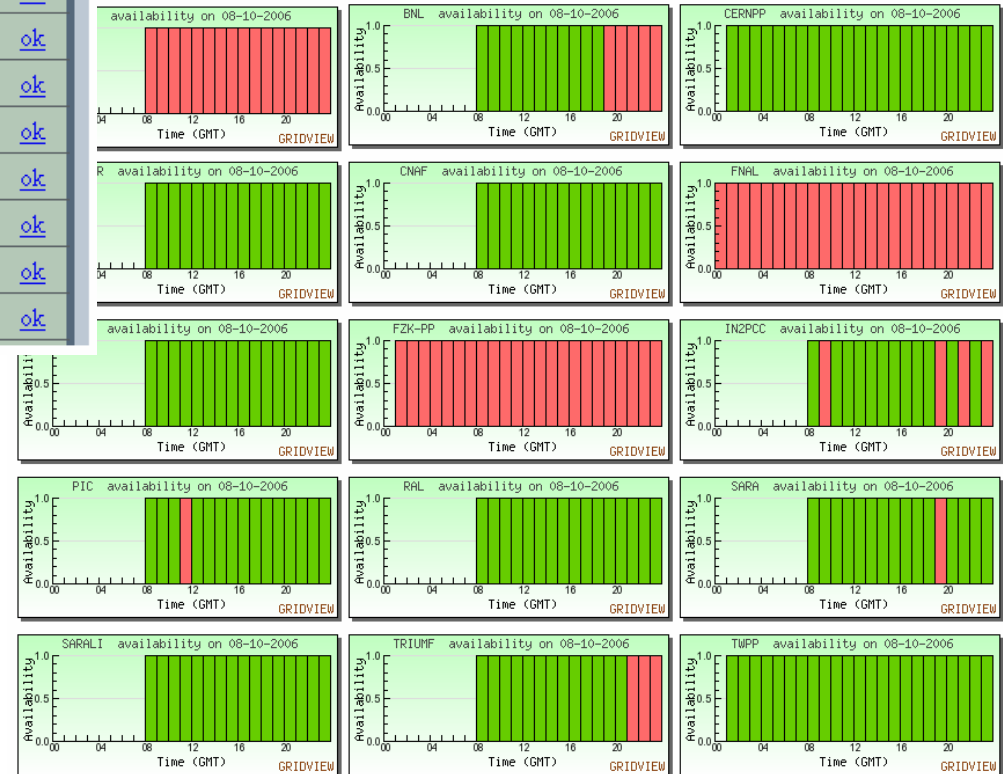
Could SAM launch a 2nd test fast after an ERROR (10min?), to discard false positives?

SAM issues: false positives

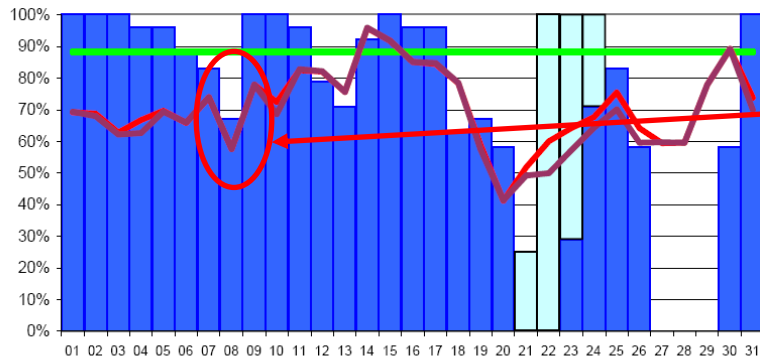
291	08-Oct-2006 18:54:36	OK	ok	3.0.2	ok	ok	ok	ok
292	08-Oct-2006 16:50:22	OK	ok	3.0.2	ok	ok	ok	ok
293	08-Oct-2006 14:53:31	OK	ok	3.0.2	ok	ok	ok	ok
294	08-Oct-2006 12:52:01	OK	ok	3.0.2	ok	ok	ok	ok
295	08-Oct-2006 10:51:19	OK	ok	3.0.2	ok	ok	ok	ok
296	08-Oct-2006 09:50:21	OK	ok	3.0.2	ok	ok	ok	ok
297	08-Oct-2006 08:50:37	OK	ok	3.0.2	ok	ok	ok	ok
298	06-Oct-2006 21:47:43	OK	ok	3.0.2	ok	ok	ok	ok
299	06-Oct-2006 19:55:53	OK	ok	3.0.2	ok	ok	ok	ok

Tier-1/0 Site Availability : Hourly Report

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



On the days 7-8 Oct some SAM data seems missing



PIC

avail: 69%

reliability: 76%

- Unavailability dip ~30% for all sites

- Could be interesting to plot as well the $MAX(availability)$ for all sites to try and detect this...

SAM issues: test rate

- IN2P3 reported about some days where too few SAM tests were run to conclude about site availability

- 21 Oct: 4 tests
- 22 Oct: 1 test !


182	23-Oct-2006 03:49:59	WARN	ok	3.0.0	warn	ok	ok	ok
183	23-Oct-2006 01:49:56	WARN	ok	3.0.0	warn	ok	ok	ok
184	22-Oct-2006 23:55:25	ERROR	ok	3.0.0	warn	ok	ok	error
185	21-Oct-2006 03:49:37	WARN	ok	3.0.0	warn	ok	ok	ok
186	21-Oct-2006 02:50:14	WARN	ok	3.0.0	warn	ok	ok	ok
187	21-Oct-2006 01:49:45	ERROR	ok	3.0.0	warn	ok	ok	error
188	21-Oct-2006 00:49:55	WARN	ok	3.0.0	warn	ok	ok	ok
189	20-Oct-2006 23:54:30	WARN	ok	3.0.0	warn	ok	ok	ok

- This was seen in all of the sites and for all of the SAM tests: Looks like a central SAM problem.

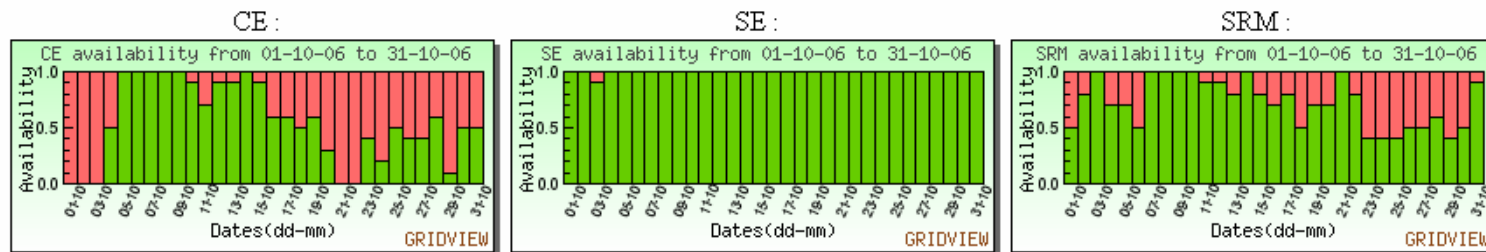
SAM issues: lcg-rm CE test

- The lcg-rm CE test introduces correlations among the CE and SRM/SE tests

Overall Service Availability for site FZK : Daily Report



Individual Service Availability for site FZK : Daily Report



- Is this what we want? This makes it more difficult to disentangle “pure CE” problems, from SE problems.

Summary

- Availability figures in the report for October look wrong for various sites (not consistent with SAM individual tests or GridView).
 - It would be good to generate a corrected report
- Most common operational issues
 - dCache scalability problem
 - CEs
 - Unavailability after urgent security patch
 - Unstable due to DoS usage pattern
- SAM framework issues/requests
 - Support for sites aggregation
 - Provide statistics for individual error results
 - Effort should be put in minimising SAM false positives
 - Overall framework problems show up in all the sites reports
 - Implement 2nd SAM fast after ERROR