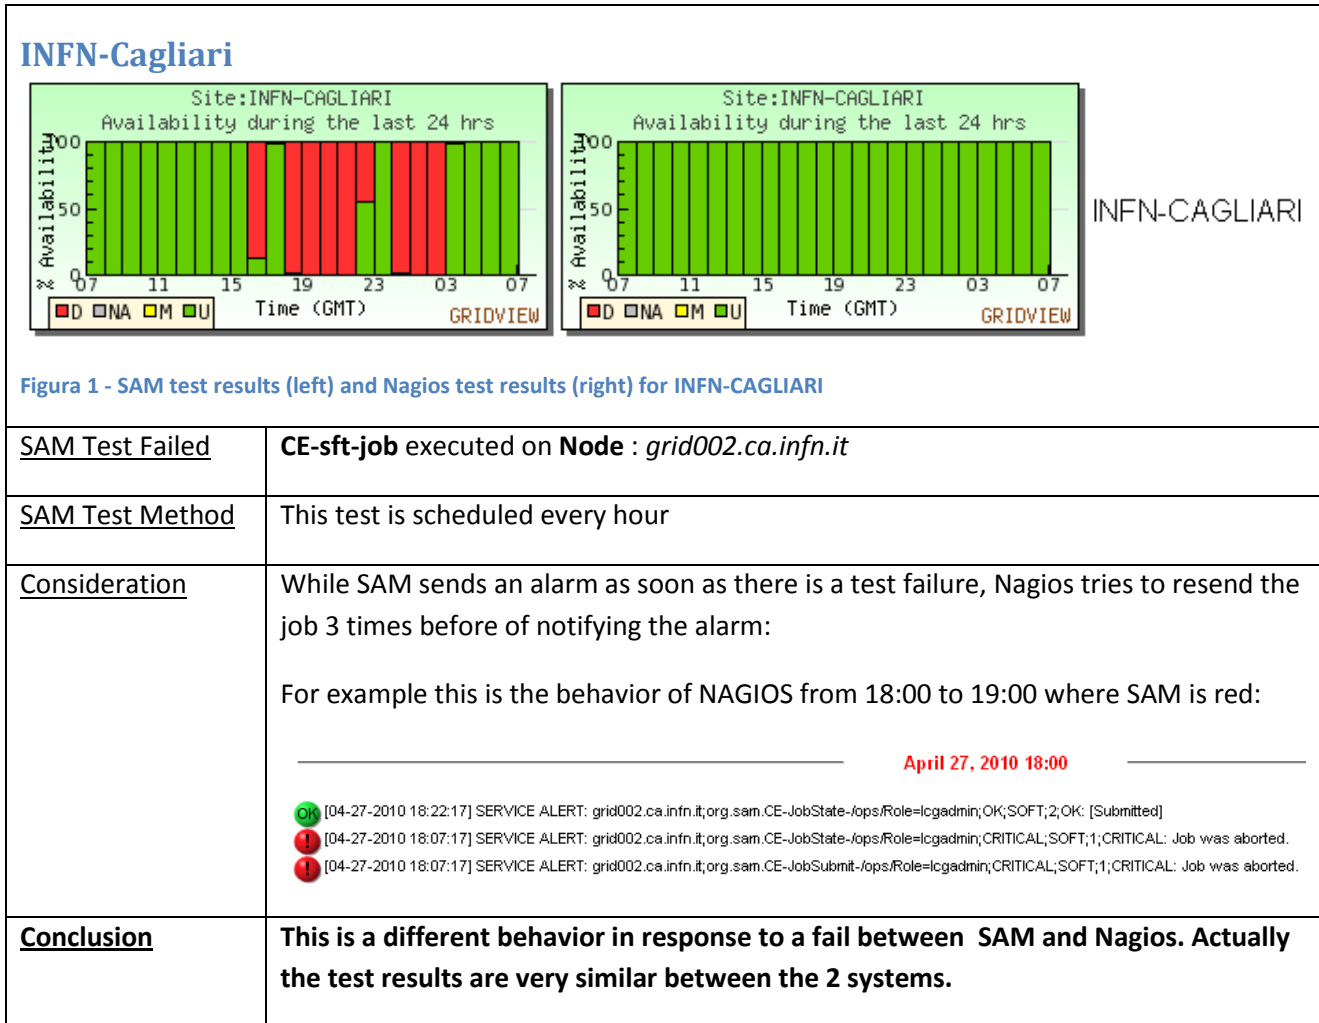


SAM and Nagios statistics comparison for IT-ROC

April 28

Comparative figures taken from <http://pps-sam.cern.ch/gridview/regions/Italy.html> , **Generated:** Wed Apr 28 03:00:22 UTC 2010 | **Region:** Italy



GRISU-COEMTA-UNIPA








Figura 3 - SAM test results (left) and Nagios test results (right) for GRISU-COMETA-UNIPA

<p><u>SAM Test Failed</u></p>	<p>CE-host-cert-valid executed on Node : unipa-ce-01.pa.pi2s2.it</p> <p>SRMv2-ls-dir executed on Node : unipa-se-01.pa.pi2s2.it</p> <p>sBDII-sanity executed on Node : unipa-ce-01.pa.pi2s2.it</p>
<p><u>Consideration</u></p>	<p>CE-host-cert-valid: it has been an isolated error so we think that the problem wasn't the certificate host</p> <p>SRMv2-ls-dir : While SAM sends an alarm as soon as there is a test failure, Nagios tries to resend the job 3 times before of notifying the alarm:</p> <p>For example this is the behavior of NAGIOS from 17:00 to 17:00 where SAM is red:</p> <p style="text-align: right;">April 27, 2010 16:00</p> <p>  [04-27-2010 16:59:05] HOST ALERT: unipa-se-01.pa.pi2s2.it;DOWN;SOFT;2;CRITICAL - Socket timeout after 10 seconds  [04-27-2010 16:57:45] HOST ALERT: unipa-se-01.pa.pi2s2.it;DOWN;SOFT;1;CRITICAL - Socket timeout after 10 seconds  [04-27-2010 16:04:23] Nagios 3.2.0 starting... (PID=9785)  [04-27-2010 16:04:23] Caught SIGHUP, restarting... </p> <p>sBDII-sanity : In nagios there isn't manifestation of this problem, probably because the problem wasn't on the host, but it was a communication problem with Gstat .</p>
<p><u>Conclusion</u></p>	<p>The test results are very similar between the 2 systems, the only real difference is due to a different behavior in response to a fail between SAM and Nagios</p>

INFN-BARI

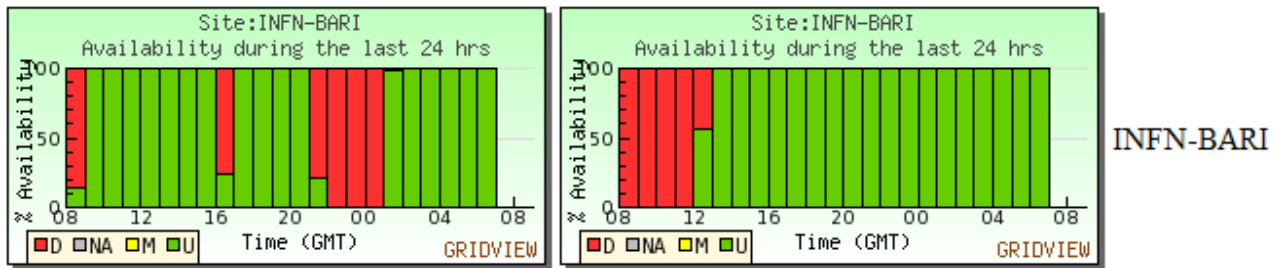


Figura 2 - SAM test results (left) and Nagios test results (right) for INFN-BARI

<u>SAM Test Failed</u>	SRMv2-put executed on Node storm-se-01.ba.infn.it
<u>SAM Failure reason</u>	[SE][StatusOfPutRequest][ETIMEDOUT] httpg://storm-se-01.ba.infn.it:8444/srm/managerv2: User timeout over lcg_cp: Connection timed out + retcode=1 + set +x Tue Apr 27 21:17:16 UTC 2010 [1272403036]
<u>Consideration</u>	<i>NAGIOS test: SRMv2-org.sam.SRM-Put.</i> Nagios doesn't send a message when the check result is tagged as SOFT, but only when the it is HARD (the status became HARD at the second consecutive failure), so the failure occurred on NAGIOS was not reported on GridView. For example, the failure (SOFT) reported at 04-27-2010 08:41:36 was not notified because the resubmission was successfully done. OK [04-27-2010 08:55:06] SERVICE ALERT: storm-se-01.ba.infn.it;org.sam.SRM-Put-ops/Role=lcgadmin;OK;SOFT;2;OK: File was copied to SRM. ! [04-27-2010 08:41:36] SERVICE ALERT: storm-se-01.ba.infn.it;org.sam.SRM-Put-ops/Role=lcgadmin;CRITICAL;SOFT;1;CRITICAL: File was NOT copied to SRM.
<u>Conclusion</u>	We have two different behavior between what SAM and NAGIOS are doing in case of failures, but the results of test are the same.

INFN-T1

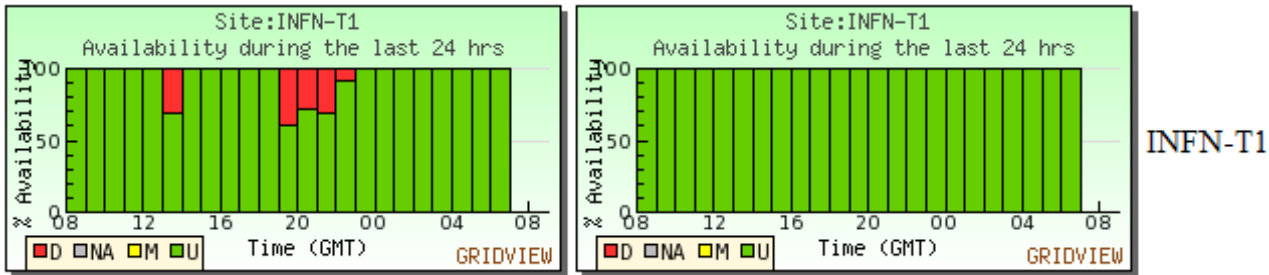



Figura 3 - SAM test results (left) and Nagios test results (right) for INFN-T1

<u>SAM Test Failed</u>	sBDII-sanity executed on Node sg01-lcg.cr.cnaf.infn.it
<u>SAM Failure reason</u>	Error message: ERROR cmd timeout 40 sec: ldapsearch -S "" -x -LLL -l 30 -h sg01-lcg.cr.cnaf.infn.it -p 2170 -b mds-vo-name=INFN-T1,o=grid * \+ Details at http://tinyurl.com/27bpxor
<u>Consideration</u>	It is not clear the behavior of the NAGIOS test <i>sBDII-org.gstat.SanityCheck</i> . At http://tinyurl.com/265ohmk it seems it is not executed every hour and there are some missing lines: at 16:26:27 on 27/04/2010 10 ok Test Env OK - errors 0, warnings 0, info 16 Detail Result at 18:26:25 on 27/04/2010 10 ok Test Env OK - errors 0, warnings 0, info 16 Detail Result at 20:26:26 on 27/04/2010 10 ok Test Env OK - errors 0, warnings 0, info 16 Detail Result
<u>Conclusion</u>	There are some missing lines in the GridView detailed test results, it is not clear if nagios do not execute the test every hour, or if there are some missing messages sent by Nagios or don't retrieved by GridView. SAM tests are not directly executed by SAM itself, but the results are taken by GSTAT.

April 26

Comparative figures taken from <http://pps-sam.cern.ch/gridview/regions/Italy.html> , **Generated:** Mon Apr 26 03:00:23 UTC 2010 | **Region:** Italy

INFN-CNAF-LHCB




INFN-CNAF-LHCB

Figura 4 - SAM test results (left) and Nagios test results (right) for INFN-CNAF-LHCB

<u>SAM Test Failed</u>	sBDII-sanity
<u>SAM Test Method</u>	This test is scheduled every 5 minutes
<u>SAM Failure reason</u>	unknown - It failed just one time, in the period under consideration, details at http://tinyurl.com/2wxod9g
<u>Consideration</u>	The result of this test is published by another service, a GSTAT server (in this case it was published from: egee017.cnaf.infn.it).
<u>Conclusion</u>	Assuming that the SAM test and the NAGIOS test are equivalent, the failure looks a temporary failure because no other tests failed in that period.

GRISU-COMETA-UNIPA

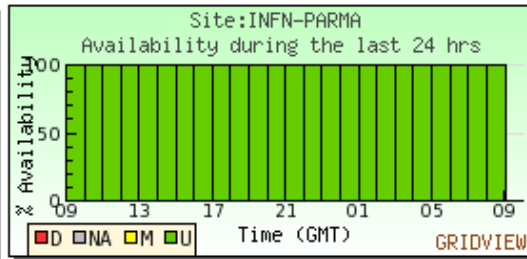
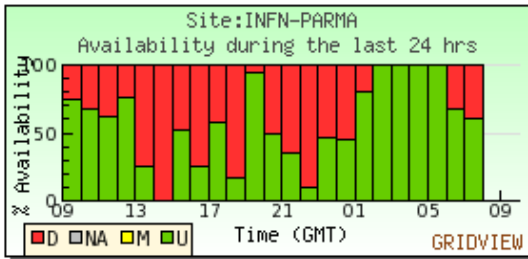


GRISU-COMETA-UNIPA

Figura 5 - SAM test results (left) and Nagios test results (right) for GRISU-COMETA-UNIPA

<u>SAM Test Failed</u>	sBDII-sanity
<u>SAM Failure reason</u>	unknown - It failed just one time, in the period under consideration, details at http://tinyurl.com/36xzd4k
<u>Consideration</u>	The result of this test is published by another service, a GSTAT server (in this case it was published from: egee017.cnaf.infn.it).
<u>Conclusion</u>	It's the same problem observed for INFN-CNAF-LHCB

INFN-Parma

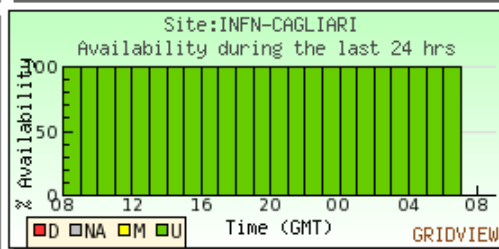
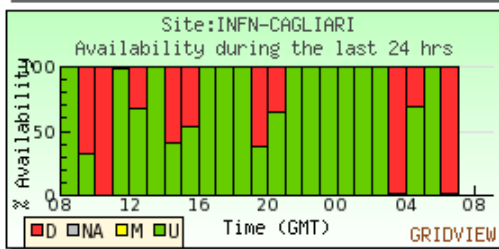


INFN-PARMA

Figura 3 - SAM test results (left) and Nagios test results (right) for INFN-Parma



The same as INFN-CNAF-LHCB

INFN-CAGLIARI



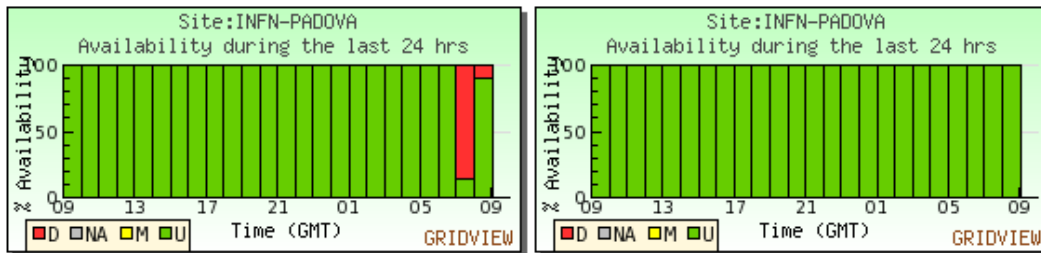
INFN-CAGLIARI

Figura 6 - SAM test results (left) and Nagios test results (right) for INFN-CAGLIARI

<u>SAM Test Failed</u>	CE-sft-job executed on Node : grid002.ca.infn.it
<u>SAM Test Method</u>	This test is scheduled every hour
<u>SAM Failure reason</u>	An example of SAM failure is available at http://tinyurl.com/2wz3c2r (the famous Condor/Maradona error message).
<u>Consideration</u>	<p>At http://tinyurl.com/2vsa3rr there are the result of Test Name : CE-org.sam.CE-JobSubmit executed and published by the nagios server samnag022.cern.ch Even if the NAGIOS test is scheduled every hour, there are not the results of all test executed.</p> <p>Looking at the Alert History view on samnag022.cern.ch for this service, we can note (as example):</p> <p style="text-align: center;">_____ April 26, 2010 _____ 01:00</p> <p> [04-26-2010 01:50:12] SERVICE ALERT: grid002.ca.infn.it;org.sam.CE-JobSubmit-/ops/Role=lgadmin;OK;SOFT;2;OK: success.</p> <p> [04-26-2010 01:05:06] SERVICE ALERT: grid002.ca.infn.it;org.sam.CE-JobSubmit-/ops/Role=lgadmin;CRITICAL;SOFT;1;CRITICAL: Job was aborted.</p>



	Nagios doesn't send a message when the check result is tagged as SOFT, but only when the it is HARD (the status became HARD at the second consecutive failure), so the failure occurred on NAGIOS was not reported on GridView.
Conclusion	So we have two different behavior between what SAM and NAGIOS are doing in case of failures, but the results of test are the same.

INFN-PADOVA



INFN-PADOVA

Figura 7 - SAM test results (left) and Nagios test results (right) for INFN-PADOVA

<u>SAM Test Failed</u>	CE-sft-lcg-rm executed on Node : prod-ce-02.pd.infn.it
<u>SAM Test Method</u>	This test is scheduled every hour
<u>SAM Failure reason</u>	An example of SAM failure is available http://tinyurl.com/2ccqzpy (CE-sft-lcg-rm-cr failure).
<u>Consideration</u>	<p>At http://tinyurl.com/269jnen there are the results of Test Name : org.sam.WN-Rep executed and published by the nagios server samnag022.cern.ch.</p> <p>Looking at the Alert History view on samnag022.cern.ch for this service, we can note (as example):</p> <p>_____ April 26, 2010 07:00 _____</p> <p> [04-26-2010 07:31:15] SERVICE ALERT: prod-ce-02.pd.infn.it;org.sam.WN-Rep-/ops/Role=lcgadmin;OK;SOFT;2;prod-wn-033.pn.pd.infn.it: OK: success.</p> <p>_____ April 26, 2010 06:00 _____</p> <p> [04-26-2010 06:31:06] SERVICE ALERT: prod-ce-02.pd.infn.it;org.sam.WN-Rep-/ops/Role=lcgadmin;CRITICAL;SOFT;1;prod-wn-017.pn.pd.infn.it: CRITICAL: METRIC FAILED [org.sam.WN-RepRep-/ops/Role=lcgadmin]: CRITICAL: File was NOT replicated to SE samdpm002.cern.ch. [ErrDB:[('lcg_util', 'server', 'CRITICAL')]]</p> <p>Nagios doesn't send a message when the check result is tagged as SOFT, but only when the it is HARD (the status became HARD at the second consecutive failure), so the failure occurred on NAGIOS was not reported on GridView.</p>

Conclusion	As observed for INFN-CAGLIARI, we have two different behavior between what SAM and NAGIOS are doing in case of failures, but the results of test are the same.
-------------------	--

TRIGRID-INFN-CATANIA




Figura 8 - SAM test results (left) and Nagios test results (right) for TRIGRID-INFN-CATANIA

SAM Test Failed	sBDII test (http://tinyurl.com/2wrb5xi)
SAM Failure reason	Unknown. "Sorry, there is no data available for the duration." Reported at http://tinyurl.com/37cqyfx
Consideration	It is not clear because SAM test can't take the results of sBDII checks executed by GSTAT. SAM retrieved sBDII results querying another service (GSTAT), while NAGIOS checks are directly executed by NAGIOS server.
Conclusion	Assuming that the SAM test and the NAGIOS test are equivalent, the unknown status reported by SAM seems not to be related to a real failure or problem on the site.

ESA-ESRIN

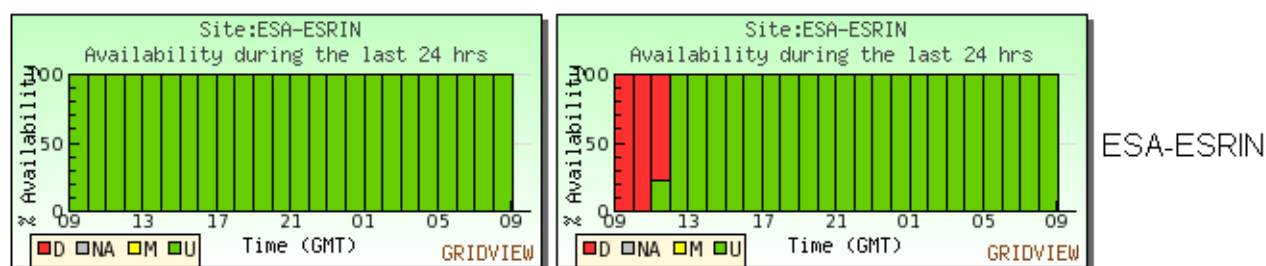



Figura 9 - SAM test results (left) and Nagios test results (right) for ESA-ESRIN

NAGIOS Test Failed	CE-org.sam.WN-Rep executed on Node : grid-eo-engine04.esrin.esa.int
NAGIOS Test Method	This test is scheduled every hour
Consideration	A t 02:00 on 25/05 the last Nagios test CE-org.sam.WN-Rep failed : [04-25-2010 02:21:04] SERVICE ALERT: grid-eo-

	<p>engine04.esrin.esa.int;org.sam.WN-Rep- /ops/Role=lcgadmin;CRITICAL;HARD;2;wnode0404.esrin.esa.int: CRITICAL: METRIC FAILED [org.sam.WN-RepCr-/ops/Role=lcgadmin]: CRITICAL: File was NOT copied to SE grid-eo-engine04.esrin.esa.int and registered in LFC prod-lfc-shared-central.cern.ch. [ErrDB:[('lcg_util', 'server', 'CRITICAL')]] CLI</p> <p>At 05:00 on 25/04 the test became successful :</p> <p> [04-25-2010 05:31:09] SERVICE ALERT: grid-eo-engine04.esrin.esa.int;org.sam.WN-Rep- /ops/Role=lcgadmin;OK;HARD;2;wnode0403.esrin.esa.int: OK: success.</p> <p>There haven't been other state change after this, Nagios sent the msg status ok at 04-25-2010 05:11:53 but Gridview report this new state only at 11:46:07</p>
<p><u>Conclusion</u></p>	<p>It seems a synchronization problem between Nagios server and GridView through msg system.</p>