

Common section:

3: # OLD Sites Functional Tests (SAM) can still be found at <https://lcg-sam.cern.ch:8443/sam/sam.py>

# For the older version of GIIS Monitoring (Gstat): <http://goc.grid.sinica.edu.tw/gstat/>

(Changed reference from SAM to NAGIOS) 6.4.5 Sites that fail NAGIOS tests but still continue to be operational

Operators need to set the ticket to `unsolvable` or set a long term deadline. A mail must be sent to the C-COD mailing list and a note put in dashboard `site notepad`. These sites should not be candidates for suspension.

## 8.1 Monitoring Tools

Removing the SAM link references.

\* Service Availability Monitoring (SAM) Tests. <https://lcg-sam.cern.ch:8443/sam/sam.py> ;

### 8.1.1 Service Availability Monitoring (SAM)

#### WHOLE SECTION NEEDS REWRITING

One of the most detailed information pages about possible problems is the Service Availability Monitoring portal (SAM Portal <https://lcg-sam.cern.ch:8443/sam/sam.py>). On the front page it gives a selection of monitored services (CE, SE, RB, etc.), regions (multiple selection list), VOs (single selection list) and the sorting order of results. For example; CE, UKI Region and OPS VO might be selected.

After the choice is made, it displays the results of the node testing scripts which are executed on a central SAM submission UI against each node of selected service type on every site. Except for CE and gCE services, a subset of test scripts is submitted to the site as a grid job and executed on one of the worker nodes. However the result is still bound to the CE to which the test job was submitted.

The test scripts contain a number of grid functional tests. As a result, a large table called the result matrix is produced. Each row of the table represents a single site and each column gives the information about the result of a single test performed against a node (or on a worker node) at a given site. Some of the tests are considered to be critical for grid operations and some of them are not.

Any failure observed in this table should be recognised, understood, diagnosed and the follow-up process should be started. This is especially important for failures of the critical tests, which should be treated with the highest priority.

The tests are executed against all sites automatically every hour. However, during the interaction with a given site ROC, the need for test resubmission may arise. In that case, the operator should manually resubmit the test job to the selected sites (on-demand resubmission). The SAM Admin web application SAMAP should be used for this purpose. (See the section further below for more information.)

The results table contains summary information about the test results (mostly OK/ERROR values). However, it is possible to see a fragment of a detailed report which corresponds to the test result by clicking on a particular result. Additionally, there are links to the GOCDB (Site Name column) and the site nodename. The detailed reports should be studied to diagnose the problems.

When investigating a particular site clicking the node name takes you to a page showing the results of the tests against that site for the last seven days. The format is a bar chart created by GRIDVIEW, clicking on a bar for

a particular test will display more detailed results. (NOTE: This view is a result of a recent update of the UI, to show the SAM history for a given site. The older style view can still be accessed by selecting the "(old history pages available here)" at the top of the Gridview History graphs area.)

In some cases, the same functionality fails on many sites at the same time. If this happens it is useful to look at the sites history. If sites that have been stable for a significant amount of time (weeks) start to fail, it might be a problem with one of the central services used to conduct the tests. Some failover procedures shall be put in place in the very near future. Since the regions have developed their own specific set of tests, it may be very helpful to contact the ROCs during the process of problem identification.

By default, the SAM Portal will show only tests that were chosen by the VO to be critical (in FCR<sup>2</sup>) and only a FCR authorised person from the VO can influence the criticality of tests. However, it is possible to configure the list of displayed tests and show non-critical tests per service type by using the multiple selection list "Tests Displayed". To revert any changes in the list of displayed tests please check the "show critical tests" check box and click the "ShowSensorTests" button. All the remaining changes (service type, region, VO) can currently only be done from the front page.

#### 8.1.1.1 Critical tests list definition (THIS NEEDS A COMPLETE REWRITE!)

Grid services are tested by specific NAGIOS tests - one test per service. Each test consists of a number of tests, some or all of which can be marked as critical. Making a test critical can be done for a number of reasons, namely

1. to give the sensor a status, thus allowing the SAM portal to apply colour-coding and display history
2. to generate an alarm in the Operations Portal<sup>2</sup> in the case that the test returns an error status
3. to include the status of the test in site availability calculations

The status of a sensor is defined as being that of the worst of its critical tests. If no tests are marked as critical, the convention is that the sensor does not have a status. Thus, in practice, at least one of a sensor's constituent tests should be defined as critical.

The SAM portal will only display the test history if the sensor that it is part of has a status (i.e. at least one test is defined as critical). This is also a pre-condition for colour-coding the display of test results. Hence the usefulness of making some tests critical.

Critical tests for the OPS VO typically generate alarms in the operations dashboard, but this must be explicitly enabled by the SAM team. Tests that will generate SAM alarms require coordination beforehand, and are the ones of greatest interest with respect to Operational Procedures.

To see which set of tests are critical for each sensor, one can use Freedom of Choice for Resources (FCR).

<https://lcg-fcr.cern.ch:8443/fcr/fcr.cgi><sup>2</sup>

On the FCR main page it is enough to select one or several services, the OPS VO (or the VO of interest), and to tick the box "show Critical Test set". It is important to leave the site as the third criteria for the grouping order (this happens by default). The resulting page will contain a list of all the tests that are critical for that VO, relative to the chosen services.

Not all of the SAM sensors are used for site availability calculations. Currently only CE, SRMv2 and sBDII sensors are used for this purpose. For those three sensors, each critical test has a direct impact on availability calculations; changes to the set of critical tests must be agreed to by various management boards, and the procedures outlined below to define a test as critical for the OPS VO must be scrupulously followed.

More information can be obtained on the SAM wiki pages dedicated to the SAM critical tests for Availability

Metrics calculations

<https://twiki.cern.ch/twiki/bin/view/LCG/EGEEWLCGCriticalProbes>

Note that SAM has both a Validation and Production infrastructure. These procedures only apply to the Production infrastructure, because the Validation instance is purely for SAM testing and does not generate SAM alarms on the ROD dashboard.

Please note that in the scope of this document we are only discussing OPS critical tests. The OPS VO tests are the basic infrastructure availability tests. Other VOs (e.g. the large HEP VOs) define implement and run additional sensors, and define their own availability calculations.

"Candidate critical tests requests" i.e. new tests which raise alarms to the dashboard upon failure(s) which operators on shift should take care of" are reviewed during the quarterly ROD Forum meetings. Urgent requests can come from the ROCs directly.

In face to face meetings, regular phone meetings or by mail, Grid OPS Procedures (aka 'Pole 2') members (currently project-eu-eg-ee-sa1-cod-pole2-proc@cernNOSPAMPLEASE.ch) will:

- \* gather input from people attending the ROD Forums meeting and make recommendations to the SAM team,
- \* check that all the C-COD shifters activities and tools are ready to handle the new critical test.

That comes down to handling the following check-list:

1. Gather C-COD feedback on the test implementation
2. Gather C-COD feedback on the impact of the general activity
3. Record feedback and final recommendations in a GGUS ticket referenced in the wiki
4. Check availability of SAM automatic masking mechanism update
5. Check availability of model mail in the dashboard for critical test to come
6. Check availability of appropriate links in the dashboard for the coming critical test
7. Check that appropriate instructions are defined in the Operations Procedure Manual - incl. In the release of the Operations Procedure Manual.

After completion of the above steps from Pole 2 in conjunction with the SAM team, the SAM team takes the final request to make a test critical to the ROC managers. Once ROC managers validate the new test to become critical for the OPS VO, it is announced to all the sites/ROCs/C-CODs via a BROADCAST during Week (W-2), announced at the Weekly Operations Meetings one week later (W-1); according to all the feedback gathered in the process, these tests will be become critical on week W at the Weekly Operations Meetings time.

The above procedure is summarized in the following list:

1. Request coming from VO/ROC/Operations
2. Evaluation and integration by the SAM team for any objection and ROC managers are informed of the new test.
3. Test put into action as a non-critical test with final "Error", "Warn" or "Ok" flags.
4. Report is submitted by VO/ROC/Operations to ROC managers including details of sites failing by region.
5. ROC reviews results with its sites.
6. Standing item on ROC managers meeting to accept, reject or request more time. A target of two weeks should be put on this step.
7. ROC managers accept, Weekly Operations Meeting notified. Wait one week.
8. Broadcast to all sites with a one week count down, mention at Weekly Operations Meeting.
9. Broadcast to all sites with one day count down. 10. Test goes critical.

### 8.1.1.2 Automatic generation of alarms in SAM DB

Every time a new test result is entered in the SAM DB, the system checks if an alarm has to be triggered or updated. This process is divided in 2 steps:

#### 1. Creation of an Alarm

The procedure inserts a new alarm in the Alarm table ONLY if:

- \* the test result is ERROR, CRIT or WARN
- \* the node belongs to a site which is Certified
- \* the VO is 'OPS'
- \* the Service name is not 'ARCCE'
- \* the test is critical for OPS VO
- \* there is no alarm already for that test, vo and node
- \* the node is not in maintenance,

At this point, the alarm is created.

Note: Only one single alarm is created for each site if APEL is failing.

#### 2. Alarm Masking

The algorithm checks if any existing alarm has to be masked due to the new alarm. If there is one or more alarms with status='new' for the OPS VO, then, for the same node and with a testid that makes it get masked according to the alarmMaskTestDef table, they are masked.

Note: alarmMaskTestDef defines which tests failing should mask other tests.

It should also be noted that SAM does not distinguish between production and pre-production nodes. When sites are "certified" and "in production" in the GOCDB, then SAM submits tests to all nodes. Also, test results are, at the time of writing, handled similarly whether submitted by SAM or SAMAP. Consequently, some pre-production nodes with monitoring status set to "on" can appear in the dashboard.

### 8.1.2 SAMAP

Section rewrite required

The SAMAP [↗](#) tool, available from the dashboard (as a separate tab) lets the operator send on demand SAM tests. Not all test sensors are available and you will only be able to send a SAM test for a site in your region.

- \* Select your region, if more than one region is shown, in the "Show Regions" list.
- \* Select the sensor from the "Node type" list.
- \* Choose other options where applicable. You may need to specifically select the WMS.
- \* Click on the "Apply Filter" button.
- \* Select the site (or sites) from the list.
- \* Click on the "Submit SAM jobs" button.

The results of the SAM test will be available after some time - this can be anything up to a few hours.

- \* Click on the "Refresh SAM jobs status" button to watch the progress of the submitted job(s).
- \* When the Current status equals DONE, then you can publish the results by clicking on the "Publish available SAM jobs results" button.

SAMAP is useful when the Site Administrator has worked on the problem, and you would like to get a quick check of the status.

## 8.1.2 NAP: NAGIOS Admin Page

THIS TOOL IS CURRENTLY UNDER DEVELOPMENT

Section rewrite required

The SAMAP<sup>?</sup> tool, available from the dashboard (as a separate tab) lets the operator send **on demand** NAGIOS tests. Not all test sensors are available and you will only be able to send a NAGIOS test for a site in your region.

- .

The results of the NAGIOS test will be available after some time - *this can be anything up to a few hours.*

## 5.1.6 Removing resources

When a *resource or node* (not site) is to be removed from the GOCDB, it's also important that it is no longer published in the site-BDII. Since the SAM DB collects both static (GOCDB) and dynamic (BDII) information, the Nagios tests will recognize the node as 'deleted', only after it disappeared from **both** the GOCDB and the Information System.

(note, the following two lines have been deleted from ROCs and Sites in this section.)

When a resource is deleted in the GOCDB, it will be flagged in the GOCDB so that the SAM DB is aware about node decommissioning. Then except for MYPROXY, VOMS, RGMA and APEL services, the node is not tested anymore.

For MYPROXY, VOMS, RGMA and APEL, there is a 1 day long retention period during which the service is still tested.

NOTE: The following REALLY should be included in the ROD manual ....

### 6.1.1.2 Automatic generation of alarms in SAM DB

-- VeraHansper - 09-Mar-2010

---

This topic: EGEE > OLDreferencesToSAM

Topic revision: r5 - 2010-04-12 - VeraHANSPERExternal



Copyright &© by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Ask a support question or Send feedback