

“NGI Services Site” for NGI A/R calculation

Problem statement

Currently it is not possible to generate NGI availability and reliability (A/R) statistics that only include “core” middleware services that are operated directly by the NGI, as typically these are hosted by different sites, and current A/R reports are per-site (i.e. only include services that are hosted by the same site in GOCDB terms). A “core” service is defined to be a middleware service that is used to access Resource Centre grid services. Examples of core services are FTS, top-BDII, VOMS, WMS and LB, etc.

An easy mechanism to circumvent the problem, is to register all NGI core middleware services in a single “NGI Services Site”. Note that core services can be operated by Resource Centres to support specific local projects. These core service instances are out of scope of this proposal. This proposal specifically applies to NGI services.

The proposed approach has the benefit that the existing infrastructure for monitoring and the computation of A/R reports can be used. It also allows the generation of NGI summarized availability statistics.

The proposed solution addresses Use Case 1 detailed in:
https://wiki.egi.eu/wiki/POEM_and_ACE_requirements

STEP 1. New GOCDB “NGI Site”

A new “NGI site” is created in GOCDB. A possible naming scheme is: “NGI_<countrycode>_SERVICES”.

See an example in GOCDB: [NGI_IT](#):

Groups

Type	Group	Actions
PRODUCTION_STATUS	Production	View
Country	IT	View
ROC	NGI_IT	View

Service Endpoints

Hostname	Service Type	Url	Description	Production	Actions
egee-bdii.cnaf.infn.it	Top-BDII	egee-bdii.cnaf.infn.it	NGI_IT top-bdii service	Y	View
ifcserver.cnaf.infn.it	Central-LFC	ifcserver.cnaf.infn.it	VO: ams02.cern.ch, argo, babar, bio, cdf, compassit, compchem, comput-er.it, cyclops, enea, enmr.eu, eticsproject.eu, euchina, euindia, glast.org, gridit, inaf, infngrid, ingv, libi, pacs.infn.it, pamela, planck, superbvo.org, theophys, tps.infn.it, virgo	Y	View
wms-lhcb.grid.cnaf.infn.it	WMS	wms-lhcb.grid.cnaf.infn.it	WMS cluster	Y	View
wms-multi.grid.cnaf.infn.it	WMS		WMS multi purpose VO cluster (pointing several WMSes instances)	Y	View

Downtimes

Users

Name	Role	E-mail	Telephone	Actions
Paolo Veronesi	Site Administrator	paolo.veronesi@cnaf.infn.it	+39 051 6092722	View

GOCDB is developed by Rutherford Appleton Laboratory, STFC, UK on behalf of EGI.eu. Licensed under the [gLite Software \(Apache\) License](#).

STEP 2: Add NGI Core services

The exact list of services that belong to a NGI site depend on which services are operated by the NGI. Examples of services are

- **TOP BDII:** egee-bdii.cnaf.infn.it – **alias** of 5 geographically distributed VMs
- **LFC:** lfcserver.cnaf.infn.it – **alias** of 2 geographically distributed physical instances
- **WMS:** wms-lhcb.grid.cnaf.infn.it – **alias** of n (n dynamically variable) geographically distributed physical instances for the LHCb VO
- **WMS:** wms-multi.grid.cnaf.infn.it – **alias** of n (n dynamically variable) geographically distributed physical instances for various regional VOs

Note 1. Only alias of distributed physical instances were added to the NGI IT site. These are monitored in SAM. Within the cluster, not working instances are automatically removed from the aliases through Nagios, so by including the alias the A/R of the cluster is considered (not of the single instance). By including the aliases the A/R of the instances of a given service are OR-ed, for example down times of a single instance will not be reflected in the A/R of the alias.

Note 2. Physical (or VMs) individual service instances that are physically hosted by different GOODB sites, were NOT REMOVED from their original hosting sites. These can be left in GOODB, as they can coexist with the alias. Having the individual instances registered can be useful in order to have alarms for the physical instances (these notifications are important even if the outage of one instance does not mean and overall service outage).

Note 3. VOMS servers are currently not added. NGI has all the VOMS servers replicated in a master/slave configuration, this means:

- No alias available (to be investigated if it's possible to have an alias for voms server, currently failover is handled on the client side)
- Web admin interface available only on the master instance
- Down of the master instance only affect users registration, proxies can still be created on the replica – need to understand if this is considered critical and should be reported in the NGI_IT site availability

So the VOMS instances even if replicated, are not totally functionally equivalent. Nevertheless, the VOMS servers can still be declared in the NGI Services site. Service flavour status can be computed by OR-ing the service endpoint status (as it is currently done for CEs, SEs etc.). The implication of this is that if VOMS 1 supports VO1 and VOMS 2 supports VO2, then if VO1 is DOWN and VO2 is UP, then:

- VOMS overall OPS Availability is 100%
- The availability experienced by VO1 is 0%
- The availability experienced by VO2 is 100%

STEP 3: Add Contact Information

The NGI Administrators staff was added as contacts for the NGI site.

STEP 4: Installation of site-BDII

In order to have the site monitored it is sufficient to create it in GOCDB , i.e. a site bdii **IS NOT NEEDED** for monitoring purposes but it was initially installed to have those service aliases published into the information system.

Now we are testing the NGI_IT site without the site bdii, monitoring and a/r calculation seem to work, but by doing so aliases are not published for the time being into the information systems.

If a site BDII is not installed and services/alias are published in other GOCDB sites, an inconsistency is introduced. The consequences of this inconsistencies were not investigated.

STATUS of NGI Grid Services in OPS tool web interfaces and PIs

GOCDB EndPoint

Site NGI_IT is available in the GOCDB programmatic interface:

https://goc.egi.eu/gocdbpi/public/?method=get_service_endpoint&sitename=NGI_IT

Example of top-BDII Status (ROC_OPERATORS profile)

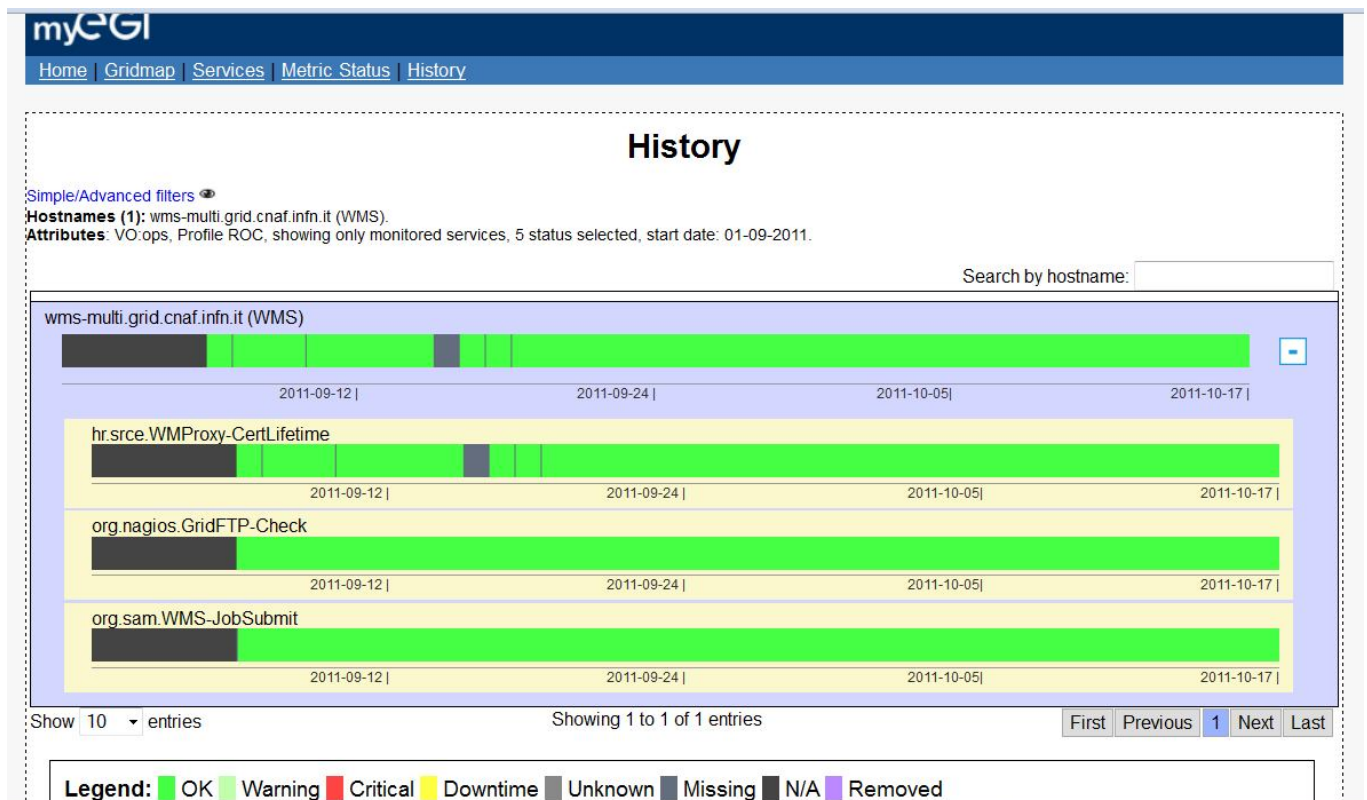
https://mon-it.cnaf.infn.it/myegi/history/?facelist_values_regions=&facelist_values_Sites=&facelist_values_services=2448%2C&vo=37&profile=2&monitored=2&status=1&status=2&status=3&status=4&status=5&startdate=01-09-2011&enddate=

The screenshot shows the 'myEGI' interface with a navigation bar containing 'Home', 'Gridmap', 'Services', 'Metric Status', and 'History'. The main content area is titled 'History' and includes a search bar for hostnames. Below the search bar, there are two horizontal bar charts. The first chart is for 'egee-bdii.cnaf.infn.it (Top-BDII)' and the second is for 'org.bdii.Entries'. Both charts show a purple segment from 2011-09-12 to 2011-09-24, followed by a green segment from 2011-09-24 to 2011-10-17. A legend at the bottom identifies the colors: green for OK, light green for Warning, red for Critical, yellow for Downtime, grey for Unknown, dark grey for Missing, black for N/A, and purple for Removed. The page also shows 'Showing 1 to 1 of 1 entries' and navigation buttons for 'First', 'Previous', '1', 'Next', and 'Last'. At the bottom, there are links for 'Help', 'Tests', 'Test Web Services', 'Support', 'Change log', and 'Web Services', along with the license 'Apache 2.0' and copyright 'CERN'.

Example of WMS Status (ROC profile)

ROC profile gives the status history of various metrics: hr.srce.WMProxy-CertLifetime, org.nagios.GridFTP-Check, org.sam.WMS-JobSubmit.

https://mon-it.cnaf.infn.it/myegi/history/?facelist_values_regions=&facelist_values_Sites=&facelist_values_services=6545%2C&vo=37&profile=5&monitored=2&status=1&status=2&status=3&status=4&status=5&startdate=01-09-2011&enddate=



Availability plots (ROC_OPERATORS)

Availability seems to be correctly computed for single instances and aggregated according to the profile chosen: ROC_OPERATORS profile:

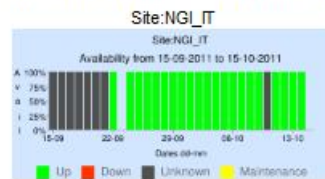
http://grid-monitoring.cern.ch/myegi/sa/?view=3&facelist_values_groups=&vo=37&profile=2&site=114164&graph=1&startdate=&enddate=

The ROC_OPERATORS profile (http://grid-monitoring.cern.ch/myegi/sam-pi/metrics_in_profiles?vo_name=ops&profile_name=ROC_OPERATORS) was chosen in order to include all the NGI core services. Currently resource centre A/R reports only include a subset of services (site-BDII, CE and SE). ROC_OPERATORS includes also LFC and FTS, WMS and VOMS metrics.

Simple/Advanced filters

Overall Service Availability for site:NGI_IT Profile:ROC_OPERATORS (Daily Report)

Availability from 15-09-2011 to 15-10-2011



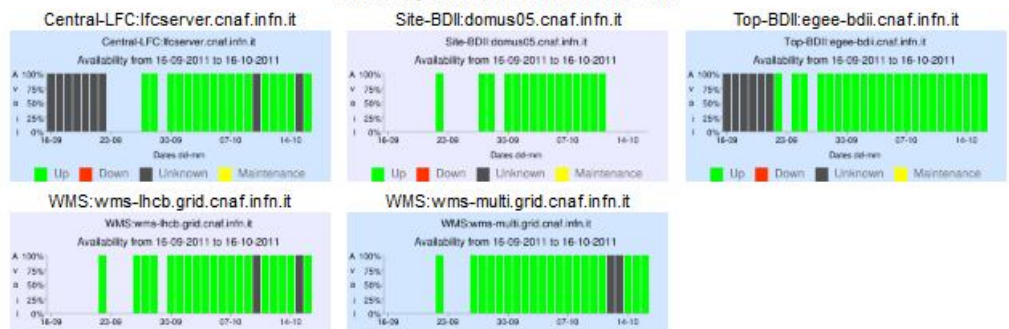
Individual Service Availability for Site:NGI_IT Profile:ROC_OPERATORS (Daily Report)

Availability from 15-09-2011 to 15-10-2011



Service Instance Availability for Site:NGI_IT Profile:ROC_OPERATORS (Daily Report)

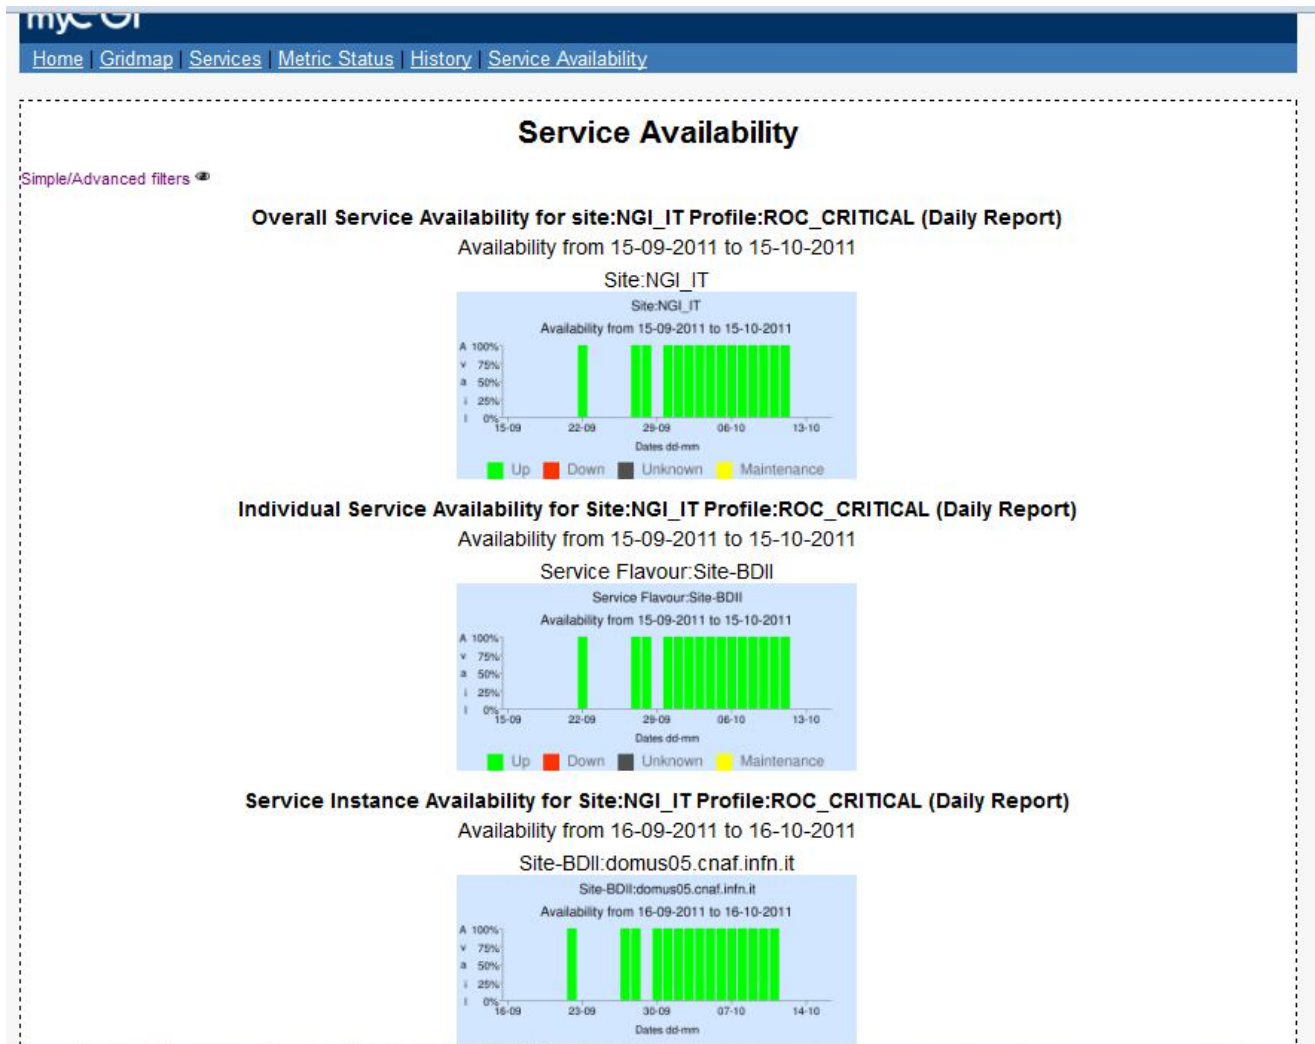
Availability from 16-09-2011 to 16-10-2011



Other profiles

ROC_CRITICAL (only site-BDII)

http://grid-monitoring.cern.ch/myegi/sa/?view=3&facelist_values_groups=&vo=37&profile=7&site=114164&graph=1&startdate=&enddate=



Availability in myEGI PI

The NGI_IT site available and reliability raw data are available in the myEGI PI – as for all standard GOCDB sites - in:

http://grid-monitoring.cern.ch/myegi/sam-pi/group_availability_in_profile?vo_name=ops&profile_name=ROC_OPERATORS&start_time=2011-10-01T00:00:00Z&end_time=2011-10-02T00:00:00Z&type=HOURLY&group_type=Site

The finest available granularity is 1 hour .

```

- <Group name="NGI_IT" type="Site">
  <Availability timestamp="2011-10-01T00:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T01:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T02:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T03:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T04:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T05:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T06:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T07:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T08:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T09:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T10:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T11:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T12:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T13:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T14:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T15:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T16:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T17:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T18:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T19:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T20:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T21:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T22:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-01T23:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
  <Availability timestamp="2011-10-02T00:00:00Z" availability="1.0" reliability="1.0" maintenance="0.0"/>
</Group>
- <Group name="UB-LCG2" type="Site">
  <Availability timestamp="2011-10-01T00:00:00Z" availability="-1.0" reliability="-1.0" maintenance="-1.0"/>
  <Availability timestamp="2011-10-01T01:00:00Z" availability="-1.0" reliability="-1.0" maintenance="-1.0"/>
  <Availability timestamp="2011-10-01T02:00:00Z" availability="-1.0" reliability="-1.0" maintenance="-1.0"/>
  <Availability timestamp="2011-10-01T03:00:00Z" availability="-1.0" reliability="-1.0" maintenance="-1.0"/>
  <Availability timestamp="2011-10-01T04:00:00Z" availability="-1.0" reliability="-1.0" maintenance="-1.0"/>

```

NGI Site in Monthly A/R Statistics

The NGI_IT site appears in September status reports – Note: It was created on the 22nd of Sept.

INFN-T1	2,252	8,192	85,516	100 %	100 %	2 %	100 %	99 %	98 %
INFN-TORINO	80	392	3,630	91 %	91 %	3 %	99 %	99 %	83 %
INFN-TRIESTE	106	382	3,300	87 %	87 %	1 %	94 %	91 %	95 %
NGI_IT	N/A	N/A	N/A	80 %	80 %	2 %	N/A	N/A	N/A
SNS-PISA	260	2,300	80,000	83 %	73 %	0 %	99 %	99 %	97 %
TRIGRID-INFN-CATANIA	160	320	2,240	100 %	100 %	3 %	96 %	93 %	85 %
UNI-PERUGIA	36	36	272	93 %	93 %	2 %	84 %	98 %	86 %

Operations Dashboard and Alarms Workflow

The NGI_IT site appears in the ops portal dashboard and we expect the same behaviour of all other sites for what concerns alarms workflow. No failures during the considered period (from 22nd of Sept till now)

Nodes

Hostname	Production	Monitored	Description	IP address
egee-bdii.cnaf.infn.it	●	●	Top-BDII	
lfcserver.cnaf.infn.it	●	●	Central-LFC	
wms-lhcb.grid.cnaf.infn.it	●	●	WMS	
wms-multi.grid.cnaf.infn.it	●	●	WMS	

Notepad : last 5 entries

Add an entry [View all entries](#) / [View only last 5 entries](#)

History

Simple/Advanced filters

Hostnames (5): domus05.cnaf.infn.it (Site-BDII), egee-bdii.cnaf.infn.it (Top-BDII), lfcserver.cnaf.infn.it (Central-LFC), wms-lhcb.grid.cnaf.infn.it (WMS), wms-multi.grid.cnaf.infn.it (WMS)

Attributes: VO:ops, Profile ROC_OPERATORS, showing only monitored services, 5 status selected, start date: 22-09-2011.

Search by hostname:

egee-bdii.cnaf.infn.it (Top-BDII)				+
Sep 28, 09:05	Oct 04, 18:11	Oct 11, 03:17	Oct 17, 12:22	
lfcserver.cnaf.infn.it (Central-LFC)				+
Sep 28, 09:05	Oct 04, 18:11	Oct 11, 03:17	Oct 17, 12:22	
wms-lhcb.grid.cnaf.infn.it (WMS)				+
Sep 28, 09:05	Oct 04, 18:11	Oct 11, 03:17	Oct 17, 12:22	
wms-multi.grid.cnaf.infn.it (WMS)				+
Sep 28, 09:05	Oct 04, 18:11	Oct 11, 03:17	Oct 17, 12:22	

Show 10 entries Showing 1 to 4 of 4 entries First Previous 1 Next Last

Legend: ■ OK ■ Warning ■ Critical ■ Downtime ■ Unknown ■ Missing ■ N/A ■ Removed

Conclusions

- The mechanism of aggregating NGI core services using GOCDB through the creation of a dedicated sites seems to work for what concerns monitoring and a/r calculation.
- NGI_IT tested it adding only aliases to distributed instances of the considered services (top bdii, lfc server, wms). Since not working instances are automatically removed from the aliases, this has the effect of ORing the test results of the physical instances. Downs of physical instances do not affect the a/r of aliases.
- N different instances of non-equivalent end-points (such as VOMS) of the same service type can be part of NGI_X_SERVICES, but the OPS A/R figures won't reflect the A/R experienced by VOs if one of the end-points is faulty (assuming the VO is supported by a subset of the end-points in the pool).

- NGI_X_SERVICES physical end-points can still be part of physical sites since it can be useful to have alarms on their status for the site admins.
- In order to have the aliases published in the information system a site bdii must be installed. But this is not mandatory for the site to appear in monitoring tools).
- ROC_OPERATORS profile can be used to get per service flavour and per NGI_X_SERVICES site A/R figures. Data can be extracted from MyEGI to produce reports.