

GridKa/KIT Service Incident Report 2010-01-13 (All times in CET)

Type of Incident: site BDII not fully functional and missing information on lcg-CEs in site BDII

Location: GridKa/KIT (FZK-LCG2)

Duration: 26 hours (including 8.5 hours scheduled downtime)

Date: January 13th, 09:30 (CET) to January 14th, 11:30 (CET)

Author: Angela Poschlad

Description

Two problems:

- 1) One of three site BDIIs in RoundRobin was not accessible from outside
Duration: 22.5 hours including 8 hours scheduled downtime
(13.01. 10:00 – 14.01. 8:30)
- 2) Missing information for lcg-CEs in information system – No matching with WMS possible
Duration: 26 hours including 8.5 hours scheduled downtime
(13.01. 09:30 – 14.01. 11:30)

Impact

- 1) Information on GridKa services missing after every third query in TopLevel BDIIs
- 2) No job submission via WMS to lcg-CEs was possible; direct submission to the CEes succeeded.

Timeline

Involved downtime of whole site 13th of January, 08:00-18:00 (CET)

<https://goc.gridops.org/downtime/list?id=52105683>

- Jan 4th: 10:00: Wrong manual entry in static lcg-CE information in Idif file
- Jan 4th – 13th: site BDII did not update cached information on lcg-CEs as wrong information was refused
- Jan 13th: 08:00: Start of scheduled downtime
09:30: reboot of site BDIIs (after update) cleared the cached CE info; missing lcg-CE information prevent WMSs to match
10:00: new site BDII in production (entry in RoundRobin)
- 18:00: After downtime SAM tests still failed but all local tests ok. Administration assumes this to be spillover effect from downtime
- 21:40 – 23:30: SAM still red -> OCE got nagios alarm; starting investigation and informing expert concentrating on published queue information but nothing found; site BDII logs empty
- Jan 14th: 07:00 – 08:30: Checking site BDII from outside; Port 2170 not reachable -> Checking firewall - ok analyzing network configuration on site BDII -> wrong routing defined on machine – sBDII fixed
- 09:00: lcg-CEs still failing but published data looked ok
- 10:40: Question raised on LCG-Rollout for WMS matching requirements gives right hint where exactly to look. Raising log level from ERROR to INFO
Finding wrong entry in Idif files from 4th of January
- 11:30: lcg-CEs fixed

Analysis

The problems were first not seen since during the scheduled downtime SAM failures were expected as normal and the local monitoring system did not show any failures. First investigations hours after the downtime was finished focused on changes made during the downtime. The problems, however, were not related to it. The log level of the site BDII's turned out to be insufficient.

The test of the site BDII from an external node finally proved a relation to the network which then could be fixed fast. The lcg-CE matching problems were still taken as a side product of the site BDII problems since nothing was changed on the CEs during the downtime. After raising the site BDII's log level the wrong manual entries in the CE's static Idif files were found.

Conclusion

Site BDII log level will stay on WARNING. Even during a downtime SAM test results will be analyzed more carefully.