

GridKa/KIT Service Incident Report 2010-06-12 (all times in CET)

Type of Incident: dCache server down

Location: GridKa/KIT

Duration: 3:15 hours

Date: Saturday, June 12, 15:30 to June 12, 18:45

Author: Doris Ressmann, Xavier Mol

Description:

A server hosting central dCache services was in kernel panic state and had to be rebooted.

Impact

Complete failure of the dCache storage element for CMS.

Timeline

- June 12: 15:38: first alarm sent to the mobile phone of the technician on-call.
- June 12: 16:26 Technician on-call started investigation of the problem via remote access.
- June 12: 18:30: Technician restarted the failing machine manually at the GridKa computing centre.
- June 12: 18:45 dCache services up and running.

Analysis

None of the available logfiles (including the bmc hardware log) state any critical issues that could have caused this failure. After the reboot, the system has been working without any detectable issues. Hence we assume that some kernel problem led to a kernel panic which could be solved only by initiating a reboot on-site.

Conclusion

This was the first time we observed this type of machine with the particular OS version in kernel panic state. We therefore assume this to be a very rare event. We will consider to provide remote access on console level ("reset button") to technicians on-call, allowing for quicker problem solving in similar cases.