

BATCH Incident at CC-IN2P3 Thursday, October 15th

Description

The parameter specifying the maximum job duration on a worker node has been unfortunately modified to a very low-level value.

The consequence is that once the long jobs finished only very short jobs were allowed to run. Long time CPU jobs still queued during the incident.

Time line of the incident

14-Oct-2009 18:30 - Farm configuration parameter has been unfortunately modified.

14-Oct-2009 22:30 - The number of running jobs is slightly low.

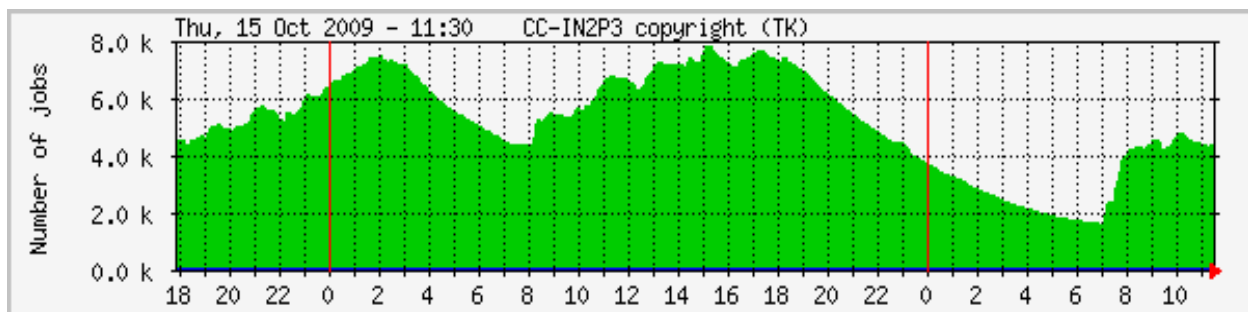
15-Oct-2009 00:30 - Minimum job running is reached. Long jobs are finished. Only a few short jobs could run.

CRITICAL alert from NAGIOS probe checking the number of running jobs.

15-Oct-2009 07:30 - Problem is noticed. The correct parameter is set up.

Batch is full available

Analysis



Follow up

- We have to create a probe to check the critical parameter values of our scheduler to be set up in our local monitoring system NAGIOS.