

# Development of a noSQL storage solution for the Panda Monitoring System

*Tuesday 7 June 2011 11:30 (30 minutes)*

For the past few years, Panda Workload Management System has been the mainstay of computing power for ATLAS experiment at the LHC. Since the start of data taking, Panda usage gradually ramped up to 840,000 jobs processed daily in the Fall of 2010, and remains at consistently high levels ever since. Given the upward trend in workload and associated monitoring data volume, the Panda team is facing a new set of challenges in the areas of database scalability and efficiency of its monitoring system. These challenges are being met with a R&D effort aimed at implementing a scalable and efficient monitoring data storage based on a noSQL solution (Cassandra). We present our motivations for using this technology, as well as data design and the techniques for efficient indexing of the specific data, which have been tested in two different hardware configurations.

**Author:** Dr POTEKHIN, Maxim (Brookhaven National Laboratory (BNL))

**Presenter:** Dr POTEKHIN, Maxim (Brookhaven National Laboratory (BNL))

**Session Classification:** Implementations II

**Track Classification:** Implementations