



Contribution ID: 115

Type: oral presentation

Studies of Big Data meta-data segmentation between relational and non-relational databases

Monday 13 April 2015 15:00 (15 minutes)

In recent years the concepts of Big Data became well established in IT-technologies. Most systems (for example Distributed Data Management or Workload Management systems) produce metadata that describes actions performed on jobs, stored data or other entities and its volume takes one to the realms of Big Data on many occasions. This metadata can be used to obtain information about the current system state, the aggregation of data for summary purposes and for statistical and trend analysis of the processes this system drives. The latter requires metadata to be stored for a long period of time. On the example of PanDA (Workload Management System for distributed production and analysis for the ATLAS experiment at the LHC and astroparticle experiments AMS and LSST) it can be seen that the growth rate of the volume of stored information has increased significantly over the last few years: from 500k completed jobs per day in 2011 up to 2 million nowadays.

Database is the central component of the PanDA architecture. Currently RDBMS (Oracle or MySQL) is used as the storage backend. To provide better performance and scalability the data stored in relational storage is partitioned into actual ("live") and archive (historical) parts. But even in this scheme, as the "archived" data volume grows, the underlying software and hardware stack encounters certain limits that negatively affect processing speed and the possibilities of metadata analysis. We had investigated a new class of database technologies commonly referred to as NoSQL databases. We suggest to use NoSQL solution for finalized, reference portion of essentially read-only data to improve performance and scalability. We had developed and implemented a heterogeneous storage which consists of both relational and non-relational databases and provides an API for unified access to stored meta-data.

We present methods of partitioning the data between two database classes, methods for efficient storage of NoSQL backend for archived data, including the analysis of different indexing schemes based on the statistics of the most frequently used queries to the historical data. We also present a comparison between different NoSQL databases to conclude their applicability to our solution. Performance of "archived" data storage in the analytical tasks is shown in the quantitative scalability and performance test results, including testing NoSQL storage against Oracle RDBMS.

This work is conducted as a part of the project to expand PanDA beyond HEP and LHC experiments.

Primary authors: Dr GRIGORIEVA, Maria (National Research Centre "Kurchatov Institute"); Ms GOLOSOVA, Marina (National Research Centre "Kurchatov Institute")

Co-authors: Dr KLIMENTOV, Alexei (Brookhaven National Laboratory); Mr RYABINKIN, Eugene (National Research Centre "Kurchatov Institute"); Mr DIMITROV, Gancho (CERN); Dr POTEKHIN, Maxim (Brookhaven National Laboratory)

Presenter: Ms GOLOSOVA, Marina (National Research Centre "Kurchatov Institute")

Session Classification: Track 3 Session

Track Classification: Track3: Data store and access