

Contribution ID: 179

Type: Sectional

Data management in heterogeneous metadata storage and access infrastructures

Friday 29 September 2017 11:50 (15 minutes)

In modern times many large projects, sooner or later, have to face the problem of how to store, manage and access huge volumes of semi-structured and loosely connected data, namely project metadata – information, required for monitoring and management of the project itself and its internal processes.

The structure of the metadata evolves all the time to meet the needs of the monitoring tasks and user requirements. And as the structure and volume of the metadata grow, it becomes impractical to store everything in a single central storage – with time such a storage becomes less flexible in structure, and query processing slows down.

To provide structure flexibility and to keep metadata access time short enough for comfort interaction with monitoring systems, next step is to replace the single central storage with a number of task-specific storages – one for active metadata, another for the archive, yet another to store aggregated information (as a cache storage), etc. In a broad sense the combination of these storages can be described as a single hybrid (or heterogeneous) metadata storage and access infrastructure.

The main goal of this infrastructure is to provide information about the project and its internal processes in a human readable and searchable way. Among the possible components of this infrastructure can be text documents, wiki pages, databases, search interfaces to storage systems, etc. To keep all these components synchronized even in case of any – software, hardware or network – failure, there is a need of some supervising tool (or a set of tools), which is aware of the infrastructure and takes care of data consistency within it.

The usual way is to create such a supervising tool individually for each case, meaning that each part of the infrastructure takes care of itself, synchronizing data only with the direct neighbors, namely the information sources for this part. And for each case one must solve same issues of reliability, throughput, scalability and fault tolerance.

To avoid solving same issues individually for every new system operating with metadata, we started to design a unified way to develop and implement such a supervising tool. It would allow developers in every particular case implement only the case-specific modules, and rest the responsibility for common issues upon the common and ready-to-use tools.

The first premise for this work appeared in 2014-2015, when we were working on the Metadata Hybrid Storage R&D project for PanDA, the workflow management system of ATLAS experiment on LHC, in NRC "Kurchatov Institute".

In this report we will explain the motivation of the problem, describe the principal architecture designed to address it and tell about the prototype system, developed and implemented for ATLAS Data Knowledge Base, the joint R&D project of NRC KI and Tomsk Polytechnic Institute, started in 2016. Also we will discuss our technology choice for the prototype, provide the performance and scalability test results and present our plans for the future.

Summary

In modern times many large projects, sooner or later, have to face the problem of how to store, manage and access huge volumes of semi-structured and loosely connected data, namely project metadata. To provide structure flexibility and to keep metadata access time short a number of task-specific storages – so-called

heterogeneous metadata storage and access infrastructure – appears. To avoid solving same issues of reliability, throughput, scalability and fault tolerance for every single subsystem in this infrastructure, we started to design a unified way to develop and implement a supervising tools for the subsystems. In the report we will present the first results and the prototype system, developed for ATLAS Data Knowledge Base.

Author: Mrs GOLOSOVA, Marina (National Research Center "Kurchatov Institute")
Co-author: RYABINKIN, Eygene (NRC "Kurchatov Institute")
Presenter: Mrs GOLOSOVA, Marina (National Research Center "Kurchatov Institute")
Session Classification: Research Data Infrastructures@Computing for Large Scale Facilities

Track Classification: Non-relational Databases and Heterogeneous Repositories