

Integration of Oracle and Hadoop: hybrid databases affordable at scale

Monday, 10 October 2016 15:45 (15 minutes)

This work reports on the activities of integrating Oracle and Hadoop technologies for CERN database services and in particular in the development of solutions for offloading data and queries from Oracle databases into Hadoop-based systems. This is of interest to increase the scalability and reduce cost for some of our largest Oracle databases. These concepts have been applied, among others, to build offline copies of controls and logging databases, which allow reports to be run without affecting critical production and also reduces the storage cost. Other use cases include making data stored in Hadoop/Hive available from Oracle SQL, which opens the possibility for building applications that integrate data from both sources.

Secondary Keyword (Optional)

Data processing workflows and frameworks/pipelines

Primary Keyword (Mandatory)

Databases

Tertiary Keyword (Optional)

Primary authors: CANALI, Luca (CERN); KOTHURI, Prasanth (CERN); BARANOWSKI, Zbigniew (CERN)

Presenter: CANALI, Luca (CERN)

Session Classification: Track 2: Offline Computing

Track Classification: Track 2: Offline Computing