Contribution ID: **96** Type: **Paper presentation**

Query Performance Comparison of PostgreSQL vs. Neo4j. A Basic Distributed Setup on OpenStack

Friday 20 September 2024 10:30 (20 minutes)

Abstract—Among the NoSQL technologies, Neo4j is one of the most popular solutions for managing graph databases and an early adopter on transactions (contrary to other NoSQL Systems). Neo4j also provides a high-level data processing language –Cypher. Despite its popularity, there are no comprehensive studies on comparing the query performance of Neo4j relative to SQL or other NoSQL counterparts. In this paper, the TPC-H benchmark database was converted from PostgreSQL to Neo4j and a set of 110 queries was devised in SQL and then translated to Cypher. The queries were executed on a OpenStack setup for both database servers following nine scenarios, by combining three database scale factors (1GB, 5GB and 10GB) with three data distribution variants (with 3, 6 and 9 nodes). Results provide support for query performance assessment of these two big data products.

Authors: Dr PINZARU, Ciprian (UAIC); EȘANU, Codrin-Stefan (UAIC); Prof. FOTACHE, Marin (Al.I. Cuza University of Iasi); CLUCI, Marius-Iulian (UAIC); Ms TEACĂ, Nicoleta (UAIC); GASNER, Paul (Agency ARNIEC/RoEduNet, Alexandru Ioan Cuza University of Iasi)

Presenter: EȘANU, Codrin-Stefan (UAIC)

Session Classification: Grid, Cloud & High Performance Computing in Science

Track Classification: Cloud Computing and Network Virtualisation