



Contribution ID: 15

Type: **Talk**

Virtual Log-Structured Storage for High-Performance Streaming

Thursday, 2 December 2021 14:00 (30 minutes)

Over the past decade, given the higher number of data sources (e.g., Cloud applications, Internet of things) and critical business demands, Big Data transitioned from batch-oriented to real-time analytics. Stream storage systems, such as Apache Kafka, are well known for their increasing role in real-time Big Data analytics. For scalable stream data ingestion and processing, they logically split a data stream topic into multiple partitions. Stream storage systems keep multiple data stream copies to protect against data loss while implementing a stream partition as a replicated log. This architectural choice enables simplified development while trading cluster size with performance and the number of streams optimally managed. This paper introduces a shared virtual log-structured storage approach for improving the cluster throughput when multiple producers and consumers write and consume in parallel data streams. Stream partitions are associated with shared replicated virtual logs transparently to the user, effectively separating the implementation of stream partitioning (and data ordering) from data replication (and durability). We implement the virtual log technique in the KerA stream storage system. When comparing with Apache Kafka, KerA improves the cluster ingestion throughput (for replication factor three) by up to 4x when multiple producers write over hundreds of data streams. Furthermore, we present the initial results of running experiments with KerA over Infiniband and Singularity in an HPC cluster.

Student?

No

Supervisor name

Supervisor email

Primary authors: Dr MARCU, Ovidiu-Cristian (University of Luxembourg); Dr NICOLAE, Bogdan (ANL); Dr COSTAN, Alexandru (University of Rennes, INSA, Inria); Dr ANTONIU, Gabriel (University of Rennes, Inria)

Presenter: Dr MARCU, Ovidiu-Cristian (University of Luxembourg)

Session Classification: HPC Technology

Track Classification: Storage and IO