

LOGICAL CLOCKS

Hadoop Performance Records Smashed by Hops and MySQL Cluster – delivers over 1.2 million operations per second.


A next-generation distribution of Apache Hadoop, called Hops, delivers a quantum leap in size and throughput of Hadoop clusters. In work with Oracle and Spotify AB, Hops delivers over 16 times the throughput of the Hadoop Filesystem (HDFS) for a real-world Hadoop workload from Spotify. Hops' key innovation is a novel distributed architecture for managing Hadoop's metadata in MySQL Cluster, Oracle's open-source NewSQL database. The result is a more scalable, reliable, and more customizable drop-in replacement for the dominant open-source Hadoop platform.

The [Hops platform](#) is a fully-compatible Hadoop distribution and is developed by Logical Clocks AB and the Distributed Computing Research Group at SICS – Swedish ICT and KTH – Royal Institute of Technology, in Stockholm. The Hops filesystem, HopsFS, scaled to 1.2 million ops/sec on Spotify's Hadoop workload and will be presented at the Spark Summit Europe on Oct 26th 2016. Hops was scaled out by adding stateless HDFS Namenodes and database nodes in MySQL Cluster. On the same hardware, HopsFS was shown to outperform HDFS by 10%, and has lower latency from even a few concurrent clients, due to the removal of the global write lock.

What's unique to Hops is the use of a distributed database to manage its metadata.

– In Hadoop, metadata is siloed away on the heap of a single Java Virtual Machine, limiting its size to a couple of hundred of GBs. In Hops, metadata is stored in MySQL Cluster and can scale out to many tens of TBs. For our benchmarks, we used a MySQL Cluster setup with 12 database nodes with replication enabled and running on commodity hosts with two E5-2620 v3 Intel CPUs and a 10 GbE network. HopsFS' 1.2 million ops/sec produced over 10 million

ops/sec on MySQL Cluster. Of existing NewSQL databases, only MySQL Cluster could provide that kind of throughput for a complex workload of multi-partition transactions with mixed index scans and primary key operations, says Jim Dowling, CEO of Logical Clocks AB and Assoc. Prof in Distributed Systems at KTH.



MySQL Cluster has been shown to scale to handle hundreds of millions of operations per second on larger clusters with 30 database nodes, so these results for HopsFS have potential for improvement with more hardware. These results have been made possible by new features in MySQL Cluster 7.5, including improvements in index/range scan performance, internal thread handling improvements, as well as improvements to the native Java connector--ClusterJ.

A research paper detailing our results is under peer review and will be released in due course. Hops, www.hops.io, is available today as open-source Apache-licensed software on [GitHub](#) and MySQL cluster is available as open-source GPLv2-licensed software from Oracle.