

Vcinity Enables Geographically Distributed Hadoop

Local Hadoop Performance on a Global Scale

EXECUTIVE SUMMARY

Challenge:

- Evaluate the performance of Hadoop with datasets distributed across 0-10,000km
- Using TeraSort, benchmark the capabilities of Vcinity's technology (Hadoop over RDMA/IB) compared to Hadoop over TCP/IP

Solution:

Vcinity's solution leverages its fabric extension technology along with a parallel file system (e.g., Lustre or Spectrum Scale) and Hadoop cluster connector to enable geographically distributed Hadoop operation.

Results:

Vcinity's testing demonstrated viable Hadoop performance across globally distributed (with 0-100ms of RTT - distances of 0 to 10,000km) datasets over a 10Gbps WAN link.

Whether it is to target consumers using data on their buying habits, provide personalized medicine based on large genomic databases, leverage surveillance data for facial recognition or study travel patterns of passengers to optimize smart city transportation, Big Data analytics is critical for various vertical markets including media, healthcare, government, manufacturing and oil & gas. Hadoop is a key component in Big Data Analytics and faces data and distance challenges just like any other Big Data application collecting, storing and accessing large datasets that tend to be spread across the globe. Running analytics on such remote datasets requires the ability to predict and move the subset of data, perceived to be the most viable, through heavy pre-processing or simply move the entirety of data brute force to the primary processing site. The inherent difficulty of needing to predict viable subsets of data ahead of time generates delays, even "forced" results, and still requires the data to be replicated after the pre-processing delays. Hadoop's own file system (HDFS) also causes additional data replication by copying data accessed by non-Hadoop workflows, resulting in additional operational expenditure.

The Vcinity Solution




Figure 1. The Vcinity Solution

By integrating our fabric extension technology with a parallel file system, which in this case is Lustre, Vcinity provides a storage overlay enabling Hadoop to access geo-diverse data without needing its own file system (i.e., HDFS). This facilitates an efficient data platform for both Hadoop and non-Hadoop workflows resulting in reduced data replication and operational expenditure compared to traditional Hadoop deployments. Vcinity's approach also leverages data where it exists today and removes most pre-processing/data reduction requirements. Vcinity eliminates the need to move data and leverages more of the data in-place to accelerate time to realize value of data analytics. In other words, we achieve geographically distributed

Hadoop where distance is no longer the barrier resulting in accelerated speed to decision and scaling performance at the pace of data explosion.

Vcinity enables the foundation – the WAN fabric – for the globally distributed Hadoop implementation with native extension of RDMA/InfiniBand over the WAN achieving higher bandwidth utilization and CPU offload, which results in substantial savings of compute and network resources. An RDMA-based parallel file system (Lustre in this case) deployment integrated with the WAN fabric offers a uniform data platform for both metadata and object-based storage, and a unified data view through consolidated management and administration. It also supports Hadoop cluster connectors for incorporating Big Data workflows and industry-standard NAS and CIFS/SMB interfaces for seamless integration with existing infrastructure.

Results

Vcinity’s testing over 10Gbps WAN link with Round Trip Times (RTT) of 0-100ms (distance of 0 to 10,000km) demonstrates viable Hadoop performance across geographically dispersed datasets as if they were local to the cluster.

Configuration details:

- Hadoop:
 - 100GB dataset
 - 3 Node Cluster (1x Master: 12c 2.6GHz 256GB RAM, 2x Slaves: 8c 3.2Ghz 256GB RAM)
 - Apache Hadoop 2.5.2 x64
- Intel® Enterprise Edition for Lustre¹ (IEEL)
 - 1x MDS, 1x MDT/MGS
 - 2x OSS, 4x OST (8 OST’s total)
- Vcinity’s fabric extension technology with 10Gbps WAN interface
- TeraSort (popular performance benchmark) is used to measure the amount of time to sort through 100GB of randomly distributed data.

Figure 2 shows the amount of time required for Vcinity to sort through 100GB of dataset scales linearly with increasing RTT compared to Hadoop over TCP/IP. Vcinity’s approach results in no performance penalty due to distance up to 20ms of RTT and better performance – even at 0ms – than TCP/IP. As Hadoop over TCP/IP runtimes degrade exponentially, they were not measured beyond RTT of 40ms.

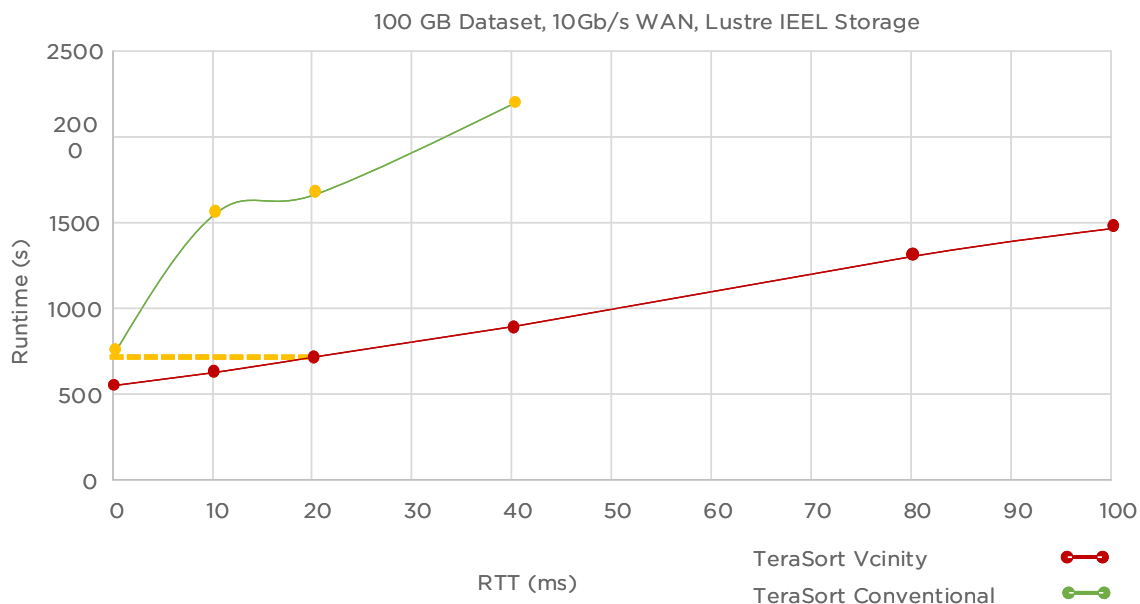


Figure 2. Vcinity’s Hadoop Testing Results

¹Trademark of Intel.

