SOFTWARE

PRODUCT BRIEF

Unstructured Data Accelerator (UDA)

Network Solutions and Hadoop Framework

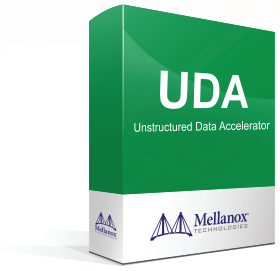

Unstructured Data Accelerator

Storing and analyzing rapidly growing amounts of data via traditional tools introduces new levels of challenges to businesses, government agencies and academic research organizations.

Hadoop MapReduce framework is a popular tool for analyzing large and unstructured data sets. Using Java based tools to process data, a data-scientist can infer users' churn pattern in retail banking, better recommend a new service to users of social media, optimize production lines based on sensor data, and detect a security breach in computer networks. Hadoop is supported by the Apache Software Foundation.

Hadoop workloads vary based on target implementation and even in the same implementation. Designing networks to sustain the different workloads introduces challenges to legacy network designs in the bandwidth and latency domains. Moving a terabyte of information can take several minutes using a 1 Giga-bit network. Minutes long operations are not acceptable in online user experiences, fraud detection and risk management tools. A better solution is required.

During MapReduce execution, data is transferred between data nodes; mapped data is delivered from the mappers to the reducer phase. At its peak, the amount of data transferred can match the amount of the incoming data. Using legacy network systems, this data transfer introduces a performance bottleneck. The Unstructured Data Accelerator (UDA) accelerates the Mapper-Reducer data transfer.

Mellanox UDA Solution

Mellanox's UDA is a user transparent software plug-in solution to the Hadoop MapReduce framework. UDA accelerates the intermediate data transfer between Mappers and Reducers.

UDA is a novel data moving protocol which uses RDMA in combination with an efficient merge-sort algorithm, to accelerate Hadoop clusters based on Mellanox InfiniBand and 10/40Gb Ethernet RoCE (RDMA over Converged Ethernet) adapter cards, to efficiently move data between data nodes in the Hadoop framework.

UDA is based on the network-levitated-merge[1] algorithm. In this algorithm the new data movement overcomes a serialization process between shuffle and merge and reduce phases. RDMA (Remote Direct Memory Access) accelerates the data transfers between mappers and reducers, as well as reducing CPU overhead by offloading data transfer burden from the CPU. Offering better CPU availability increases the number of processes available for analytics, increasing throughput capability.

UDA parallelizes shuffle and merge processes with the reduce phase, to enable UDA parallelism Map output Files (MoF) should be available and complete. The new processing scheme adds a significant performance boost to the framework by better utilizing CPU cores availability and less re-submission of jobs due to failed merge process.

UDA Performance

Unstructured Data Accelerator can double data analytics throughput and reduce total job execution time by up to 50 percent. Larger data sets will benefit from higher

HIGHLIGHTS

BENEFITS

- Integrated support for FDR InfiniBand or 10/40GbE to MapReduce Framework
- Leverage data transfer with RDMA technology and efficient merge-sort algorithm
- Reduce job execution time by up-to 50%
- Better CPU utilization for map and reduce jobs
- Seamless integration to Hadoop framework

performance boost, due to the better wire throughput and lower latency. More CPU slots enable better allocation for Mappers and Reducers jobs in the framework timeline. Fewer hard disk access requests, provided by the memory-to-memory transactions, result in faster data transactions and overall faster execution time.

UDA is a scalable solution deployed on clusters with three data nodes and clusters with thousands of data nodes. Performance results growth are near linear with the cluster size.

UDA Availability

UDA is an open-source project, developed under the Apache 2.0 license.

The most up-to-date UDA package can be downloaded from: https://code.google.com/p/uda-plugin/

UDA was jointly developed by the Parallel Architecture and System Laboratory headed by Dr. Weikuan Yu from Auburn University and Mellanox.

References

[1] "Hadoop acceleration through network levitated merge"; Yandong Wang, Xinyu Que, Weikuan Yu, Goldenberg, D., Sehgal, D., International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2011

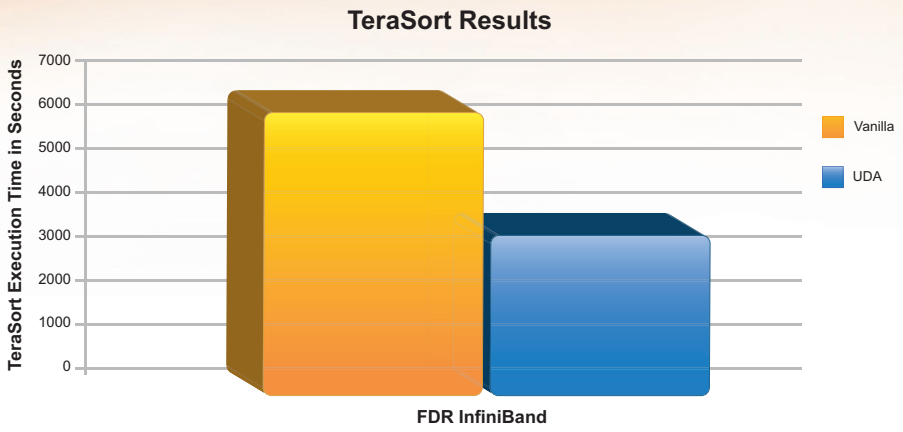

Mellanox TECHNOLOGIES

350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com

15-557PB Rev 1.0