



Delivering In-memory Computing Using Mellanox Ethernet Infrastructure and MinIO's Object Storage Solution

EXECUTIVE SUMMARY

Analytic tools such as Spark, Presto and Hive are transforming how enterprises interact with and derive value from their data. Designed to be in memory, these computing and analytical frameworks process volumes of data 100x faster than Hadoop Map/Reduce and HDFS - transforming batch processing tasks into real-time analysis. These advancements have created new business models while accelerating the process of digital transformation for existing enterprises.

A critical component in this revolution is the performance of the networking and storage infrastructure that is deployed in support of these modern computing applications. Considering the volumes of data that must be ingested, stored, and analyzed, it quickly becomes evident that the storage architecture must be both highly performant and massively scalable.

This solution brief outlines how the promise of in-memory computing can be delivered using high-speed Mellanox Ethernet infrastructure and MinIO's ultra-high performance object storage solution.

IN MEMORY COMPUTING

With data constantly flowing from multiple sources - logfiles, time series data, vehicles, sensors, and instruments - the compute infrastructure must constantly improve to analyze data in real time. In-memory computing applications, which load data into the memory of a cluster of servers thereby enabling parallel processing, are achieving speeds up to 100x faster than traditional Hadoop clusters that use MapReduce to analyze and HDFS to store data.

Although Hadoop was critical to helping enterprises understand the art of the possible in big data analytics, other applications such as Spark, Presto, Hive, H2O.ai, and Kafka have proven to be more effective and efficient tools for analyzing data. The reality of running large Hadoop clusters is one of immense complexity, requiring expensive administrators and a highly inefficient aggregation of compute and storage. This has driven the adoption of tools like Spark

KEY BUSINESS BENEFITS

MinIO and Mellanox: Better Together

High performance object storage requires the right server and networking components. With industry-leading performance combined with the best innovation to accelerate data infrastructure Mellanox provides the networking foundation needed to connect in-memory computing applications with MinIO high performance object storage. Together, they allow in-memory compute applications to access and process large amounts of data to provide high speed business insights.

Simple to Deploy, Simpler to Manage

MinIO can be installed and configured within minutes simply by downloading a single binary and executing it. The amount of configuration options and variations has been kept to a minimum resulting in near-zero system administration tasks and few paths to failures. Upgrading MinIO is done with a single command which is non-disruptive and incurs zero downtime.

MinIO is distributed under the terms of the Apache* License Version 2.0 and is actively developed on Github. MinIO's development community starts with the MinIO engineering team and includes all of the 4,500 members of MinIO's Slack Workspace. Since 2015 MinIO has gathered over 16K stars on Github making it one of the top 25 Golang* projects based on a number of stars.

which are simpler to use and take advantage of the massive benefits afforded by disaggregating storage and compute. These solutions, based on low cost, memory dense compute nodes allow developers to move analytic workloads into memory where they execute faster, thereby enabling a new class of real time, analytical use cases.

These modern applications are built using cloud-native technologies and, in turn, use cloud-native storage. The emerging standard for both the public and private cloud, object storage is prized for its near infinite scalability and simplicity - storing data in its native format while offering many of the same features as block or file. By pairing object storage with high speed, high bandwidth networking and robust compute enterprises can achieve remarkable price/performance results.

DISAGGREGATE COMPUTE AND STORAGE

Designed in an era of slow 1GbE networks, Hadoop (MapReduce and HDFS) achieved its performance by moving compute tasks closer to the data. A Hadoop cluster often consists of many 100s or 1000s of server nodes that combine both compute and storage.

The YARN scheduler first identifies where the data resides, then distributes the jobs to the specific HDFS nodes. This architecture can deliver performance, but at a high price - measured in low compute utilization, costs to manage, and costs associated with its complexity at scale. Also, in practice, enterprises don't experience high levels of data locality with the results being suboptimal performance.

Due to improvements in storage and interconnect technologies speeds it has become possible to send and receive data remotely at high speeds with little (less than 1 microsecond) to no latency difference than if the storage were local to the compute.

As a result, it is now possible to separate storage from the compute with no performance penalty. Data analysis is still possible in near real time because the interconnect between the storage and the compute is fast enough to support such demands.

By combining dense compute nodes, large amounts of RAM, ultra-high-speed networks and fast object storage, enterprises are able to disaggregate storage from compute creating the flexibility to upgrade, replace, or add individual resources independently. This also allows for better planning for future growth as compute and storage can be added independently and when necessary, improving utilization and budget control.

Multiple processing clusters can now share high performance object storage so that different types of processing, such as advanced queries, AI model training, and streaming data analysis, can run on their own independent clusters while sharing the same data stored on the object storage. The result is superior performance and vastly improved economics.

HIGH PERFORMANCE OBJECT STORAGE

With in-memory computing, it is now possible to process volumes of data much faster than with Hadoop Map/Reduce and HDFS. Supporting these applications requires a modern data infrastructure with a storage foundation that is able to provide both the performance required by these applications and the scalability to handle the immense volume of data created by the modern enterprise.


Building large clusters of storage is best done by combining simple building blocks together, an approach proven out by the hyper-scalers. By joining one cluster with many other clusters, MinIO can grow to provide a single, planet-wide global namespace. MinIO's object storage server has a wide range of optimized, enterprise-grade features including erasure code and bitrot protection for data integrity, identity management, access management, WORM and encryption for data security and continuous replication and lambda compute for dynamic, distributed data.

MinIO object storage is the only solution that provides throughput rates over 100GB/sec and scales easily to store 1000s of Petabytes of data under a single namespace. MinIO runs Spark queries faster, captures streaming data more effectively, and shortens the time needed to test, train and deploy AI algorithms.

LATENCY AND THROUGHPUT

Industry-leading performance and IT efficiency combined with the best of open innovation assist in accelerating big data analytics workloads which require intensive processing. The Mellanox ConnectX® adapters reduce the CPU overhead through advanced hardware-based stateless offloads and flow steering engines. This allows big data applications utilizing TCP or UDP over IP transport to achieve the highest throughput, allowing completion of heavier analytic workloads in less time for big data clusters so organizations can unlock and efficiently scale data-driven insights while increasing application densities for their business.

Mellanox Spectrum® Open Ethernet switches feature consistently low latency and can support a variety of non-blocking, lossless fabric designs while delivering data at line-rate speeds. Spectrum switches can be deployed in a modern spine-leaf topology to efficiently and easily scale for future needs. Spectrum also delivers packet processing without buffer fairness concerns. The single shared buffer in Mellanox switches eliminates the need to manage port mapping and greatly simplifies deployment. In an





object storage environment, fluid resource pools will greatly benefit from fair load balancing. As a result, Mellanox switches are able to deliver optimal and predictable network performance for data analytics workloads.

The Mellanox 25, 50 or 100G Ethernet adapters along with Spectrum switches results in an industry leading end-to-end, high bandwidth, low latency Ethernet fabric. The combination of in-memory processing for applications and high-performance object storage from MinIO along with reduced latency and throughput improvements made possible by Mellanox interconnects creates a modern data center infrastructure that provides a simple yet highly performant and scalable foundation for AI, ML, and Big Data workloads.

CONCLUSION

Advanced applications that use in-memory computing, such as Spark, Presto and Hive, are revealing business opportunities to act in real-time on information pulled from large volumes of data. These applications are cloud native, which means they are designed to run on the computing resources in the cloud, a place where Hadoop HDFS is being replaced in favor of using data infrastructures that disaggregate storage from compute. These applications now use object storage as the primary storage vehicle whether running in the cloud or on-premises.

Employing Mellanox networking and MinIO object storage allows enterprises to disaggregate compute from storage achieving both performance and scalability. By connecting dense processing nodes to MinIO object storage nodes with high performance Mellanox networking enterprises can deploy object storage solutions that can provide throughput rates over 100GB/sec and scales easily to store 1000s of Petabytes of data under a single namespace. The joint solution allows queries to run faster, capture streaming data more effectively, and shortens the time needed to test, train and deploy AI algorithms, effectively replacing existing Hadoop clusters with a data infrastructure solution, based on in-memory computing, that consumes a smaller data center footprint yet provides significantly more performance.

WANT TO LEARN MORE?

Click the link below to learn more about object storage from MinIO VAST: <https://min.io/>

Follow the link below to learn more about Mellanox end-to-end Ethernet storage fabric: <http://www.mellanox.com/ethernet-storage-fabric/>

About MinIO

Founded in 2014, MinIO is now the world's fastest growing object storage system. Backed by some of the smartest minds in storage and venture capital including Nexus, General Catalyst, Dell Technologies Capital, Intel Capital, AME Cloud Ventures and key angel investors, the company has raised \$23.3M through its Series A round. Learn more about MinIO at www.min.io.

About Mellanox

Mellanox Technologies is a leading supplier of end-to-end InfiniBand and Ethernet interconnect solutions and services for servers and storage. Mellanox interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance capability. Mellanox offers a choice of fast interconnect products: adapters, switches, software, cables and silicon that accelerate application runtime and maximize business results for a wide range of markets including high-performance computing, enterprise data centers, Web 2.0, cloud, storage and financial services.

To find out more, visit our website: www.mellanox.com



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