

Accelerate SQL on Hadoop for Enterprises with Cisco UCS and Dremio Solution



Highlights

Real-time performance with linear scalability

Cisco UCS® Integrated Infrastructure for Big Data and Analytics, combined with Dremio distributed architecture, provides a simplified intelligent data infrastructure with high performance and scalability to meet growing real-time business demands.

Integrated infrastructure built on Cisco UCS architecture

Cisco UCS Integrated Infrastructure for Big Data and Analytics is a proven platform for enterprise analytics applications. The latest generation of the Cisco Unified Computing System™ (Cisco UCS) empowers customers to refine increasing amounts and varieties of data into business value faster than ever before, accelerating the journey toward digital transformation.

Ease of deployment

Cisco UCS Manager simplifies infrastructure deployment with an automated, policy-based mechanism that helps reduce configuration errors and system downtime. Cisco UCS Central Software enables users to manage multiple Cisco UCS instances or domains across disparate locations and environments.

Accelerated analytics

Dremio's open-source technology allows users to discover, curate, accelerate, and share data at a high rate with a high level of concurrency for large-scale analytical application processing. The Dremio Data Reflections feature provides in-memory query optimization to accelerate query processing.

Advanced query optimization with Dremio Learning Engine

The Dremio Learning Engine suggests transformations and joins, performs predictive caching of metadata, and detects changes in schemas for physical data sets. It is designed to reduce the cost of moving data on the network and uses scatter and gather read and write operations. It has a zero-serialization and deserialization design, allowing low-cost data movement between servers.

Full integration with YARN

Dremio has a scale-out architecture. The Data Reflections store is provisioned and managed as a YARN application in big data Hadoop clusters ranging from hundreds to thousands of servers and can take advantage of the aggregate memory of a cluster.

Challenges of the existing data environment

With the rapid increase in data from various sources, analytics processing has become increasingly complex.

Business intelligence tools, data science systems, and machine-learning models works best when data resides in a single, high-performance relational database.

Today, most organizations use one or more nonrelational data stores, Hadoop, and NoSQL databases. Many newer data stores are more agile than before and provide improved scalability, but at a cost in speed and ease of access through traditional SQL-based analysis tools. Additionally, the raw data in these stores often is too complex or inconsistent for analysis with business intelligence tools. With the adoption of Hadoop, NoSQL, and other big data technologies, data analytics has become more dependent on IT than ever, and business intelligence users and data scientists are underutilized as a result.

SQL on Hadoop

A SQL query on Hadoop (typically run as a Hive job) translates to multiple jobs underneath the entire Hadoop cluster to read and process this data. The first generation of these jobs were in map-reduce and subsequent generations used tools like Spark or Tez which brought in-memory execution by greatly reducing the time for Hive SQL queries. However, the latencies for these queries were still large to enable interactive use of BI tools by business analysts. It has been a challenge to use popular BI tools directly with data on Hadoop as the SQL queries are not interactive in nature due to the humongous volume of data that is stored on the Hadoop cluster.

Unique solution from Cisco and Dremio

The Cisco Unified Computing System™ (Cisco UCS®) together with Dremio technology enables end users to run interactive queries on Hadoop. Business intelligence tools can directly connect and fetch data from Hadoop clusters through Dremio without having to stage this data or build another layer of data warehousing to drive the business intelligence tools.

Dremio addresses the problems of modern data stores by providing an experience that is easily and immediately usable by analysts, with the integration of common business intelligence tools, R, Python, and any SQL-based tool. This feature frees consumer business from potentially costly data workflows that depend on custom solutions from IT and data professionals, while simultaneously granting instant compatibility with industry-standard tools that are already deployed across millions of desktops.

Dremio accelerates end-user analysis tasks by using the Apache Arrow columnar in-memory data technology and an intelligent caching feature called the Dremio Accelerator. With acceleration enabled, SQL queries can be run interactively and are accelerated with the zero-copy and zero serialization and deserialization structure. This approach allows the same data to be used by multiple end users, which is essential to efficient use of the memory and computing resources that make the model work.

In addition, processors have evolved, becoming more sophisticated with cache locality, pipelining, and single instruction multiple data (SIMD) optimization. As a query optimizer, Dremio maintains physically optimized representations of source data (data sets) known as Data Reflections. It can accelerate a query by using one or more Data Reflections to partially or entirely satisfy that query, rather than processing the raw data in the underlying data source. Big data customers who use the Hadoop File System (HDFS) particularly can benefit from this approach from the way data is stored in Parquet and is retrieved in the in-memory Arrow columnar representation.

Dremio’s columnar query processing and acceleration help ensure interactive performance on any data volume. However, data engineers, business analysts, and data scientists need a way to curate data to make it suitable for the needs of specific projects. This requirement represents a fundamental shift from the IT-centric model, in which consumers of data initiate a request for a data set and wait for IT to fulfill that request weeks or months later. Dremio enables a self-service model, in which consumers of data use Dremio’s data curation capabilities to collaboratively discover, curate, accelerate, and share data without relying on IT. This user experience is powered by a modern, intuitive, web-based user interface.

In Hadoop YARN deployment mode, Dremio integrates with the YARN resource manager to secure computing resources in a shared multitenant environment. Figure 1 shows the high-level deployment architecture of Dremio on a Hadoop cluster. Dremio coordinators and executors are configured to use HDFS volumes for the cache and spill directories.

Figure 1. Dremio YARN deployment

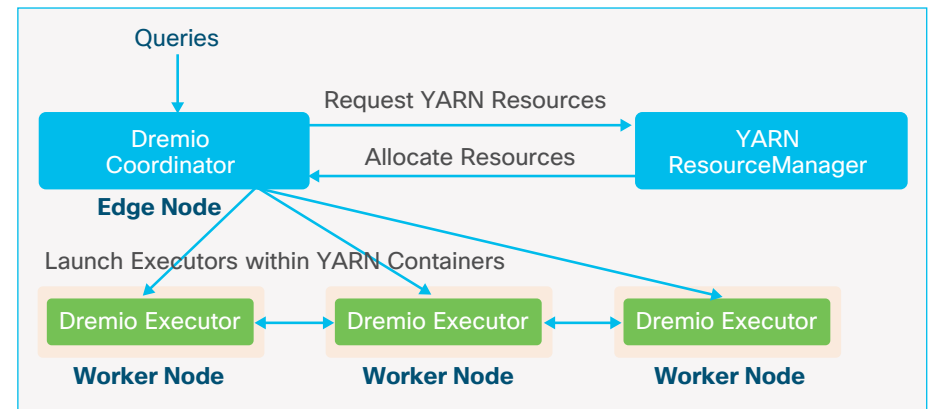
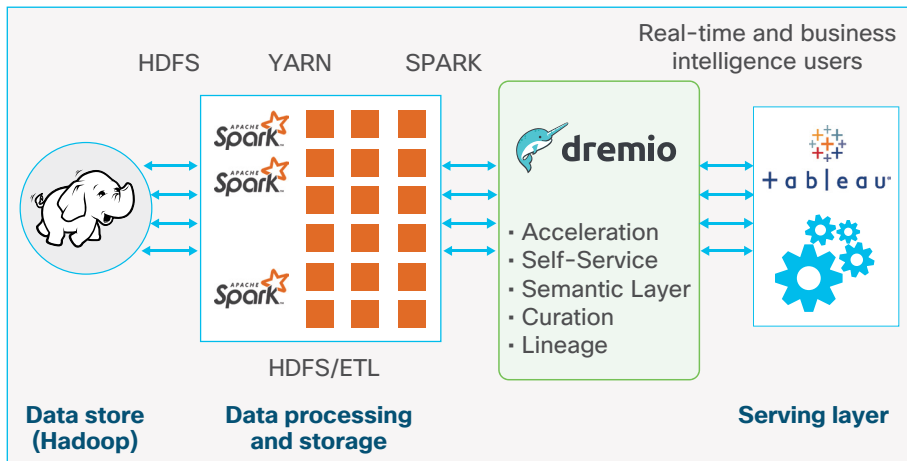


Figure 2. Dremio as a new tier for data analytics



Cisco UCS Integrated Infrastructure for Big Data and Analytics

Cisco UCS Integrated Infrastructure for Big Data and Analytics provides an end-to-end architecture for processing high volumes of real-time or archived data, both structured and unstructured. At the same time, it quickly and efficiently delivers out-of-the-box performance while scaling from small to very large deployments as business requirements and big data and analytics requirements grow.

Organizations today must be sure that the underlying physical infrastructure can be deployed, scaled, and managed in a way that is agile enough to change as workloads and business requirements change. Cisco UCS Integrated Infrastructure for Big Data and Analytics has redefined the potential of the data center with its revolutionary approach to managing computing, network, and storage resources to successfully address the business needs of IT innovation and acceleration.

Cisco UCS 6300 Series Fabric Interconnects

Cisco UCS 6300 Series Fabric Interconnects provide high-bandwidth, low-latency connectivity for servers, with Cisco UCS Manager providing integrated, unified management for all connected devices. Cisco UCS 6300 Series Fabric Interconnects are a core part of Cisco UCS, providing low-latency, lossless 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions.

Cisco® UCS Fabric Interconnects offer the full active-active redundancy, performance, and exceptional scalability needed to support the large number of nodes that are typical in clusters serving big data applications. Cisco UCS Manager enables rapid and consistent server configuration using service profiles and automates ongoing system maintenance activities such as firmware updates across the entire cluster as a single operation. Cisco UCS Manager also offers advanced monitoring with options to raise alarms and send notifications about the health of the entire cluster.

Cisco UCS C240 M5 Rack Servers

Cisco UCS C240 M5 Rack Servers are dual-socket, two-rack-unit (2RU) servers offering industry-leading performance and expandability for a wide range of storage and I/O-intensive infrastructure workloads, such as big data, analytics, and collaboration. These servers use the new Intel® Xeon® Scalable processor family with up to 28 cores per socket. They support up to 24 double-data-rate 4 (DDR4) dual in-line memory modules (DIMMs) for improved performance and lower power consumption. The DIMM slots are 3D XPoint ready, supporting next-generation nonvolatile memory technology.

Depending on the server type, Cisco UCS rack servers have a range of storage options. For big data Hadoop workloads, Cisco recommends the Cisco UCS C240 M5 Rack Server, which supports up to 24 small form-factor (SFF) 2.5-inch drives with a Cisco 12-Gbps SAS Modular RAID Controller. In addition, all the servers have two modular M.2 cards that you can use for bootup. A modular LAN-on-motherboard (mLOM) slot supports dual 40 Gigabit Ethernet network connectivity with the Cisco UCS Virtual Interface Card (VIC) 1387.

Reference architecture

Dremio can be deployed on a Hadoop cluster on Cisco UCS Integrated Infrastructure for Big Data and Analytics on a few or all of the data nodes for query processing and can be co-located on the data nodes. For the reference architecture for Cisco UCS Integrated Infrastructure for Big Data and Analytics. See <https://blogs.cisco.com/datacenter/cpav5>

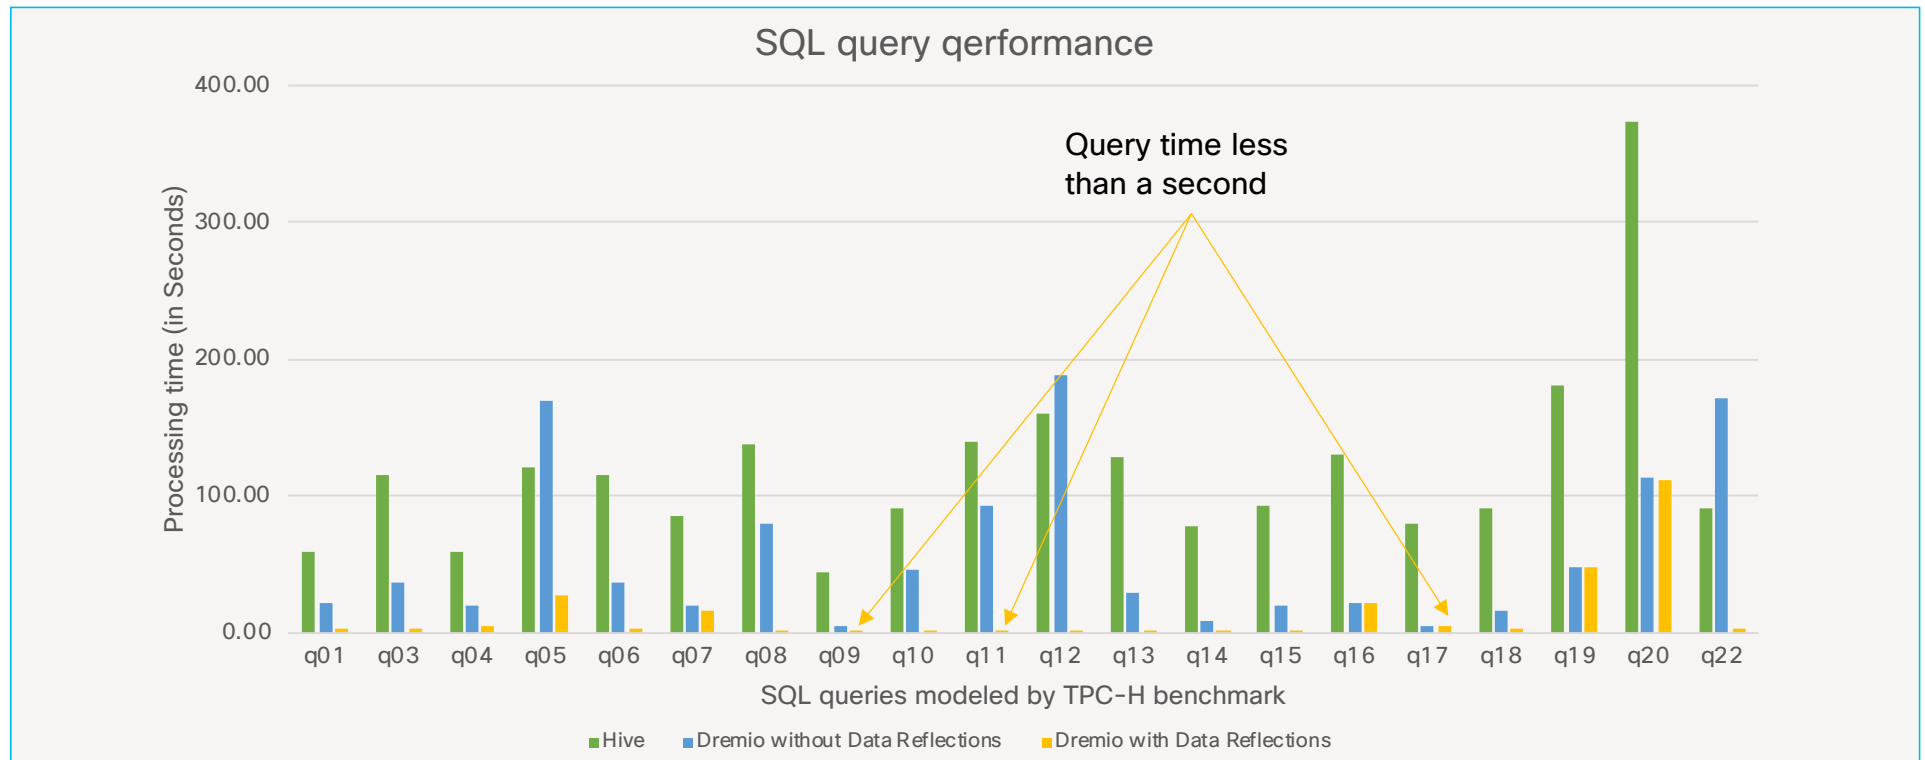
Performance benchmark and results

A benchmark modeled on the TPC-H query performance benchmark was used to test performance with 1 TB of data on Cisco UCS Integrated Infrastructure for Big Data and Analytics with Dremio. The system consisted of nine Cisco UCS C240 M5 servers with one management node and eight data nodes, and Dremio was installed on all eight data nodes. The data

node configuration consisted of twenty-four 1.8-TB 10,000-rpm SFF hard-disk drives (HDDs), two Intel Xeon Scalable 6132 processors, and 192 GB of DRAM. It was connected to dual 40-Gbps links in an active-standby configuration connected to a pair of third-generation Cisco UCS fabric interconnects.

The benchmark consisting of 22 queries were run using 1 TB scale factor with Dremio's in-memory columnar query processing (with and without Data Reflections) and with just Hive on Hadoop over data sets on HDFS. These tests were run mainly to demonstrate Dremio's query performance (with and without Data Reflections) compared to that of Hadoop HiveQL, which queried the same data sets. In addition, performance was increased by orders of magnitude (10x to 100x) for specific queries with Data Reflections and vectorized processing enabled.

Figure 3. Runtime comparison: TPC-H Apache Hive versus Dremio Data Reflections (lower is better)



Conclusion

Cisco UCS Integrated Infrastructure for Big Data and Analytics with Dremio open-source technology provides a simplified, self-service, intelligent infrastructure and a data analytics capability with the scalability needed to meet the growing business demands of customers while eliminating the cost and complexity of ETL tools.

Cisco UCS technology together with Dremio Data Reflections query-processing acceleration and Apache Arrow technology combines the power of Hadoop with a dependable YARN deployment model that can be implemented rapidly using Cisco UCS C-Series Rack Servers. This solution could potentially accelerate workloads in orders of magnitude as demonstrated in the benchmark tests.

For more information

For additional information, see the following resources:

- **Cisco UCS Integrated Infrastructure for Big Data and Analytics:** <https://blogs.cisco.com/datacenter/cpav5>
- **Cisco UCS big data solutions:** <http://www.cisco.com/go/bigdata>
- **Dremio:** <http://www.dremio.com/>

