

Action Vector in Hadoop for SQL Analytics

The Fastest, Most Industrialized SQL in Hadoop

A Technical Overview

Table of Contents

Overview	3
Action Vector in Hadoop - Uniquely Fast	6
Exploiting the CPU.....	6
Exploiting Single Instruction, Multiple Data (SIMD).....	7
Utilizing CPU Cache as Execution Memory	7
Parallel Execution.....	7
Column-based Storage.....	8
Supporting Updates via Positional Delta Trees (PDTs).....	8
Data Compression.....	8
Partitioned Table Support.....	9
Transparent Node Failover	9
YARN for Resource Management	9
Action Director for Ease of Administration	10
Conclusion	11

Action Vector in Hadoop for SQL Analytics

The Fastest, Most Industrialized SQL in Hadoop

Apache Hadoop has captured the imagination of those responsible for large-scale data management and data analytics projects. Many view it as a game-changer to the enterprise. However, it is not without challenges. It requires rare and expensive skillsets, creates long and error-prone implementation cycles, lacks support for popular reporting and BI tools, and executes with inadequate speed. This has led to a search for an alternative that combines all of the benefits of Hadoop with SQL – the world’s most widely used data querying language.

Overview

The “SQL on Hadoop” marketplace is crowded with solutions that fall into three broad categories:

- **SQL outside Hadoop:** The customer deploys both a Hadoop Cluster and a DBMS Cluster, on the same or separate hardware, and uses a connector to pass data back and forth between the two systems. The approach is expensive and hard to manage, and is most often adopted by traditional data warehouse vendors. Solutions that fall into this category include Vertica, Teradata and Oracle.
- **SQL alongside Hadoop:** These vendors have taken an existing SQL engine and modified it such that when generating a query execution plan it can determine which parts of the query should be executed via MapReduce, and which parts should be executed via SQL operators. Data that is processed via SQL operators is copied off HDFS into a local table structure. Solutions that fall into this category include Hadapt, RainStor, Citus Data and Splice Machine.
- **SQL on Hadoop:** These vendors are building SQL engines, from the ground up, that enable native SQL processing of data in HDFS while avoiding MapReduce. These products have limited SQL language support, rudimentary query optimizers, which can require handcrafting queries to specify the join order, and no support for trickle updates. Product immaturity is reflected in their lack of support for internationalization, limited security features and lack of workload management. Solutions that fall into this category include Cloudera Impala, Apache Drill, Pivotal HAWQ and IBM’s Big SQL.

A new “SQL in Hadoop” category is required for Actian’s revolutionary SQL analytics in Hadoop capability. Actian Vector in Hadoop™ (“VectorH”) represents the next generation of innovation for Hadoop and employs a unique approach to bring all of the benefits of industrialized, high performance SQL together with Hadoop as part of a comprehensive platform.

Actian VectorH was designed and built working closely with leading organizations across multiple industries to understand their big data requirements, including SQL access to Hadoop data. As a result, VectorH is not only the most comprehensive SQL in Hadoop analytics platform but also the fastest. There are several key capabilities that make Actian VectorH different from all other SQL in Hadoop approaches:

1. **Comprehensive:** A SQL in Hadoop capability shouldn't be used in isolation. Organizations need to view SQL in Hadoop as part of a complete analytics solution that covers the end-to-end process, from connecting to raw big data, analyzing data natively on Hadoop, all the way through to delivering analytics results that impact business. Actian VectorH connects to many data types and sources, then optimizes query planning and execution to deliver high-performance analytical processing and SQL access to analytical results through your favorite visualization tools—all natively on Hadoop.
2. **Trickle Updates:** The ability to not only access but also update/append Hadoop data is important. Most SQL on Hadoop solutions will claim some form of a bulk update as well as bulk appends but it's the changes to individual records or the removal of individual records that is the big challenge on HDFS. Actian VectorH is the ONLY offering in the market able to insert, update or delete individual records. It is also fully ACID-compliant, with multi-version read consistency, plus system-wide failover protection.
3. **Data Quality:** Organizations need to ensure that Hadoop data accessed by business users has been curated for high quality. This is especially true if businesses will use the Hadoop data for operational decision making. Actian VectorH supports data blending and enrichment, and quality validation to ensure business users are accessing high quality Hadoop data.
4. **SQL:** Actian VectorH supports standard ANSI SQL to make it easy and flexible to work with SQL-based applications, as well as standard business intelligence or visualization tools. This also ensures that any existing SQL queries will be able to access Hadoop data without the need for modification. In addition, the Actian database supports advanced analytics, such as cubing and window functions.
5. **Performance Optimized:** Actian VectorH is built using mature technology strongly rooted in data management with published TPC benchmarks. It includes a proven cost-based query planner and optimizer that make optimal use of all available resources from the node, memory, cache, and all the way down to the CPU.
6. **Security:** Actian VectorH includes native enterprise-grade DBMS security including authentication, user- and role-based security, data protection and encryption.
7. **Compression:** Hadoop can reduce storage costs, but it still requires replicating copies for high availability. Actian VectorH compresses data by typically a factor of 6 and stores it in a native compressed columnar format for faster SQL performance.

8. **Manageability:** Actian VectorH is certified YARN-ready, proving it uses the Hadoop YARN capability to automatically manage resources and nodes on the cluster. The platform includes a web-based database administration and querying tool to manage analytics and query processing on your Hadoop cluster. The platform also includes native Backup and Restore capabilities.
9. **Architecture:** Actian VectorH is built for extreme scalability. It easily expands and scales as the size of your data needs grow. Plus, it is architected to handle extremely complex queries, thousands of nodes and SQL users, and petabytes and more of data.
10. **Native:** Actian VectorH runs natively in Hadoop, eliminating the need to move data off the HDFS nodes into a separate database or file system. It is extremely flexible and supports the most commonly used Hadoop distributions, including Apache, Cloudera, Hortonworks, and MapR.

Actian VectorH complements other capabilities delivered within the Hadoop environment, including Spark for data integration, streaming analytics, and machine learning:

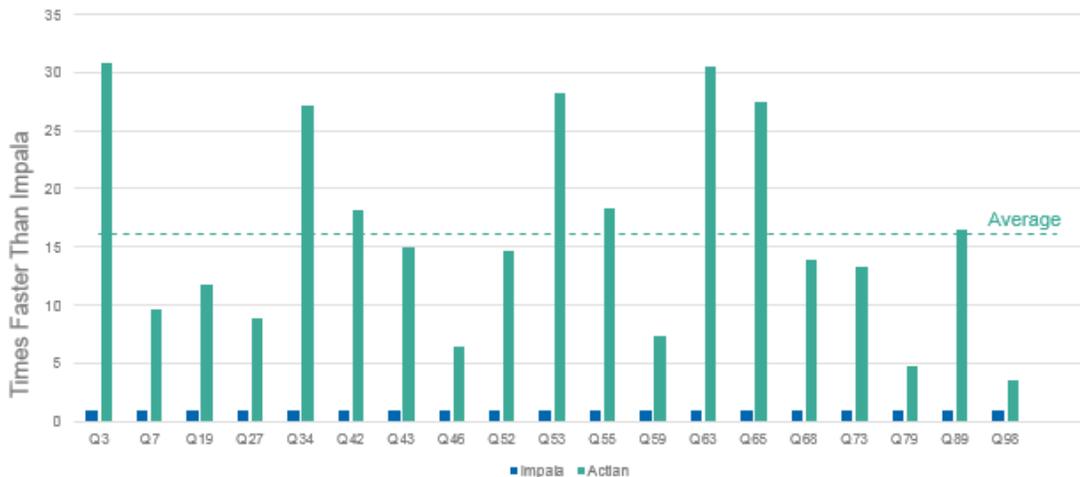
- **Actian Vector in Hadoop** – fastest columnar analytics database for native SQL in Hadoop. Actian VectorH is a mature technology, which has been hardened in the enterprise and includes support for localization and internationalization, advanced security features and workload management. Actian VectorH contains a mature RDBMS engine that performs native SQL processing of data in HDFS. It has rich SQL language support, an advanced query optimizer, support for trickle updates, and has been certified for use with the most popular BI tools. Plus, it has been benchmarked to perform more than 30 times faster than other approaches to SQL on Hadoop.

By bringing the raw performance of a proven columnar analytics database to the data nodes in the Hadoop Cluster, Actian VectorH is able to achieve performance levels never before seen with any SQL on Hadoop solution. The Actian platform can process Hadoop data natively in HDFS an order of magnitude faster than any other SQL on Hadoop solution and provides industrialized SQL access and enterprise class security to HDFS data. Fast SQL access to Hadoop data opens up new opportunities. Think not only about support for larger data sets, more users and more complex workloads, but also the ability to query Hadoop data using standard BI and Reporting tools.

This paper explains how Actian SQL analytics in Hadoop achieves extremely fast performance for typical data warehouse and data mart workloads requiring high performance interactive SQL in Hadoop.

Action Vector in Hadoop - Uniquely Fast

Action VectorH has been proven in both benchmarking tests and real-world customer implementations to be the fastest SQL in Hadoop in the market. Figure 1 shows actual benchmark results comparing VectorH to Cloudera Impala for the same queries running on equivalent hardware. On average, VectorH was about 16 times faster than Impala. But for more complex queries, VectorH was more than 30 times faster. Interestingly, because VectorH is the only SQL in Hadoop to support “updates,” it can run all 99 TPC-DS queries. Impala can only run the subset of TPC-DS queries because it lacks support for “updates.”



Both Executed on the Same Hardware and Software Environment:
5 Node Cluster with 64GB of RAM per node and 12x1TB Hard Disks

Background to "Impala Subset" of TPC-DS benchmark can be found here:
<http://blog.cloudera.com/blog/2014/01/impala-performance-dbms-class-speed/>

Figure 1 – Action Vector in Hadoop Benchmark Comparison against Cloudera Impala

To understand what makes Action VectorH extremely fast, let’s review some of the key features that make it unique.

Exploiting the CPU

Action VectorH takes advantage of powerful CPU features that other databases can’t or don’t. During the past three decades CPU processing capacity has roughly followed Moore’s Law, doubling roughly every two years. However, today the improvements in CPU data processing performance are not just the result of increases in clock speed and the number of transistors on the chip. CPU manufacturers have introduced additional performance features such as multi-core CPUs and multi-threading, which are transparently leveraged by most database software.

There are, however, other optimizations that were introduced in the last decade that are typically not transparently leveraged by most database software. Examples include so-called SIMD instructions, larger chip caches, super-scalar functions, out-of-order execution and hardware-accelerated string-based operations. In fact, most of today’s database software that was originally written in the 1970s or 1980s has become so complex that in order to take advantage of these

performance features a complete rewrite of the database software would be required.

Action Vector was written from the ground up to take advantage of performance features in modern CPUs, resulting in dramatically higher data processing rates compared to other relational databases. Action VectorH leverages these innovations and brings this unbridled processing power to the data nodes in a Hadoop Cluster.

Exploiting Single Instruction, Multiple Data (SIMD)

SIMD enables a single operation to be applied on every entity in a set of data all at once. Action VectorH takes advantage of SIMD instructions by processing vectors of data through the Streaming SIMD Extensions instruction set. Because typical data analysis queries process large volumes of data, the use of SIMD may result in the average computation against a single data value taking less than a single CPU cycle .

At the CPU level, traditional databases process data one tuple at a time, spending most of the CPU time on overhead to manage tuples and not on the actual processing. In contrast, Action VectorH processes vectors of hundreds or thousands of elements at once, which effectively eliminates these overheads. Some of the leading solutions being built to provide SQL access on Hadoop have expressed plans to leverage the vectorization techniques for which Action Vector is known.

Utilizing CPU Cache as Execution Memory

The majority of the improvements to database server memory (RAM) over the last number of years have resulted in much larger memory pools, but not necessarily faster access to memory. As a result, relative to the ever-increasing clock speed of the CPU, access to memory has become slower over time. In addition, with more CPU cores requiring access to the shared memory pool, contention can be a bottleneck to data processing performance.

In order to achieve maximum data processing performance, Action VectorH avoids the use of shared RAM as execution memory. Instead, Action VectorH uses the private CPU core and CPU caches on the Hadoop data nodes as execution memory, delivering significantly greater data processing throughput.

Parallel Execution

Action VectorH implements a flexible adaptive parallel execution algorithm and can be scaled-up or scaled-out to meet specific workload requirements. Action VectorH can execute statements in parallel using any number of CPU cores on a server or across any number of data nodes on a Hadoop cluster.

Action VectorH follows a traditional Master-Worker pattern. The Hadoop NameNode serves as the master node and the Hadoop DataNode serves as the workers. But to avoid potentially competing with Hadoop for processing resources, Action VectorH is installed on the DataNodes with fail-over capabilities if desired. Taking the raw power of the Vector data processing engine to the HDFS data is what gives Action VectorH its unique performance characteristics.

Column-based Storage

When relational database software was first written, it implemented so-called row-based storage: all data values for a row are stored together in a data block (page). Data was always retrieved row-by-row, even if a query only accessed a subset of the columns. This storage model worked well for On-Line Transaction Processing (OLTP) systems in which data was stored highly normalized, tables were relatively narrow, queries often retrieved very few rows and many small transactions were processed.

Data warehouse solutions that are built on Hadoop are different:

- Tables are often (partially) denormalized, resulting in many more columns per table, not all of which are accessed by most queries.
- Most queries retrieve many rows.
- Data is added through a controlled rather than ad-hoc process and often large data sets are added at once or through an ongoing (controlled) stream of data.

As a result of these differences, a row-based storage model typically generates unnecessary I/O for a Hadoop data warehouse workload. The Actian column-based storage model, in which data is stored together in data blocks (pages) on a column-by-column basis, is generally accepted as a superior storage model for data analysis queries.

Supporting Updates via Positional Delta Trees (PDTs)

Actian VectorH uniquely implements a fully ACID-compliant transactional database with multi-version read consistency. Any new transaction will see all previously committed transactions, both small incremental transactions and large bulk data loads. Changes are always written persistently to a transaction log before a commit completes to always ensure full recoverability.

One of the biggest challenges with HDFS is that it is not designed for incremental updates. VectorH is the only solution that provides the ability to update individual records with minimal impact on read-performance. Some vendors are attempting to copy VectorH by implementing poor solutions that keep track of the updates using disk instead of memory. These approaches are dreadfully slow, which is why they are switched off by default.

Actian VectorH uniquely addresses this challenge using our patented high-performance in-memory Positional Delta Trees (PDTs), which are used to store small incremental changes (inserts that are not appends), as well as updates and deletes.

Conceptually a PDT is an in-memory structure that stores the position and the change (delta) at that position. Queries efficiently merge the changes in PDTs with data stored on HDFS. Because of the in-memory nature of PDTs, small DML statements can be processed very efficiently. A background process writes the in-memory changes to disk once a memory threshold is exceeded.

Data Compression

The algorithms Actian VectorH uses to compress data have been selected for their speed of

decompression over a maximum compression ratio. The compression ratio you can achieve with Actian VectorH is highly data-dependent. 4-6x compression ratios are very common for real-world data, but both lower and higher compression ratios have been observed in the lab.

In order to improve I/O performance, Actian VectorH allocates a portion of physical memory for a memory-based disk buffer, the Column Buffer Manager (CBM). Data is automatically pre-fetched from disk and stored in the CBM, mirroring the data as it is stored on HDFS. In contrast to many other databases, Actian VectorH does not decompress data in the memory buffer, but rather data is decompressed only once it is ready for data processing on the DataNode.

Actian VectorH automatically chooses the most optimal compression on a page by page basis and there can be multiple different algorithms in use. Decompression comes at almost no cost because it is directly integrated in the vector-based processing. The Actian VectorH decompression is far more efficient than alternative speed-optimized compression libraries such as LZOP that many other products have utilized.

Partitioned Table Support

Actian VectorH includes table partitioning that splits a logical table into multiple physical tables on HDFS using hash-partitioning on a column (or set of columns). Table partitioning is introduced during bulk load, where each DataNode is allocated data determined by a partitioning key. The scalability witnessed in Actian VectorH can be attributed in part to elimination of all-to-all communications through the ability to execute joins locally on each DataNode.

Transparent Node Failover

In case of a DataNode failure, the responsibilities of the failed nodes have to be reassigned to the remaining nodes in the worker-set, such that the load is spread as evenly as possible with respect to data locality. This process is handled transparently by querying the HDFS NameNode to find out where the partitions are located, and using this information to determine the missing replicas and re-replicating them to other (new) nodes. In case of a master node failure, VectorH will automatically start up on a secondary master node and continue processing as though nothing happened.

YARN for Resource Management

Historically, MapReduce applications were the only applications on a Hadoop cluster that could be managed by the Hadoop resource management capabilities. In the latest versions of Hadoop, YARN provides containers for all YARN Ready applications launching on worker nodes. This helps control usage of CPU, memory, disk, and network resources. YARN provides resource negotiation and management for the entire Hadoop cluster and all applications running on it. By integrating with YARN, non-MapReduce based distributed applications, such as Actian VectorH query workloads, can run as first-class citizens on Hadoop clusters, sharing resources side-by-side with MapReduce based applications.

Actian VectorH also uniquely includes a block placement API that enables the ability to override the default HDFS block placement strategy to maximize performance. It can automatically recover from a data-node failure to rebalance the data and get the performance back up to expected levels.

Action Director for Ease of Administration

Action Director is a management console to manage the VectorH analytics database. It provides a simple point and click interface for managing, administering and querying Action VectorH instances on premise and in the cloud. Available for both Windows and Linux, Action Director has been extended to include support for the new features included in Action VectorH in Hadoop including partitioned tables and HDFS browsing.

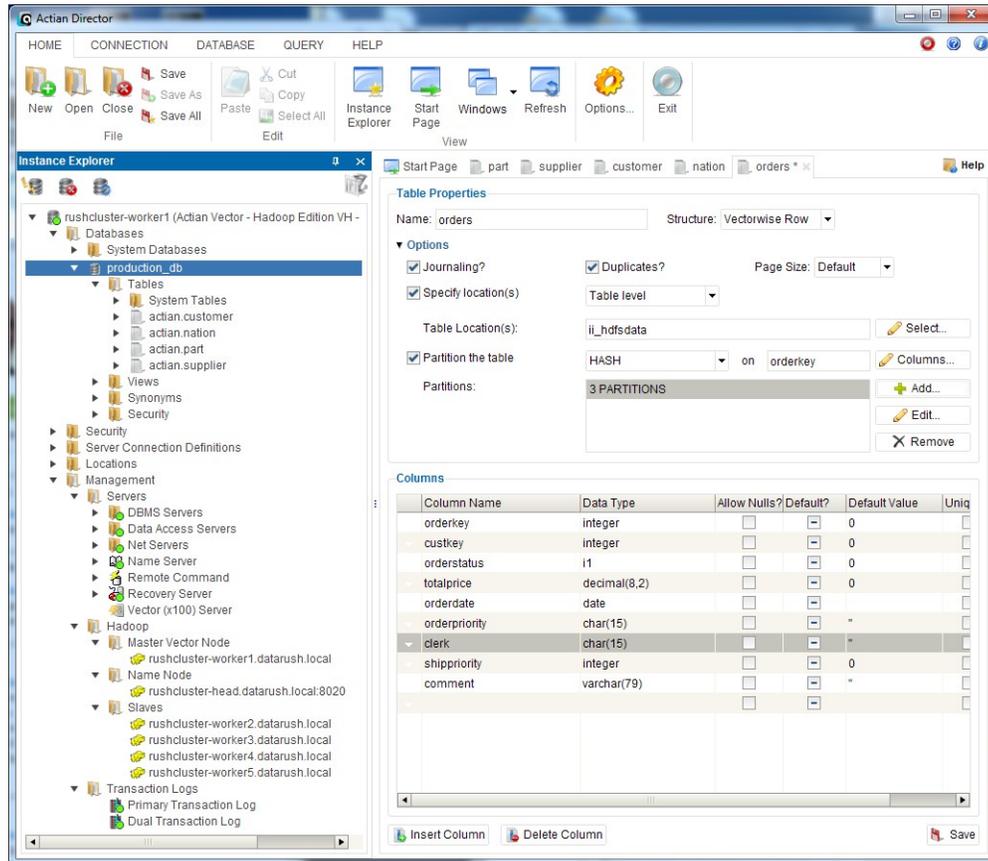


Figure 2 – Action Director Console User Interface

Conclusion

If you need to analyze large volumes of data on Hadoop and you don't want to take the risk of an expensive or lengthy implementation project, your choice should be Actian VectorH. Implement an easy to deploy, easy to use, ANSI compliant solution and benefit from significantly better query performance than any other SQL analytics on Hadoop solution. Actian VectorH is the foundation for revolutionary performance gains in database processing – gains that are so game changing that they appear on our competitors' long-term roadmaps.

Rest assured that there is more innovation to come! Future versions of Actian VectorH will not only introduce new functionality and leverage new Hadoop capabilities, but also continue to leverage CPU performance features and implement other optimizations to get absolute maximum analytic processing and query performance on Hadoop. Use Actian VectorH if you are looking for the world's fastest, most industrialized SQL analytics database running natively in Hadoop.

About Actian:

Actian transforms big data into business value for organizations of all sizes, no matter where they are on their analytics journey. We help companies win by empowering them to connect to data of any type, size or location; analyze it quickly wherever it resides; and take immediate action on accurate insights gained to delight their customers, gain competitive advantage, manage risk and find new sources of income. With the Actian Vector in Hadoop, Actian has delivered the world's fastest analytics database built to run 100 percent natively in Hadoop. Among tens of thousands of organizations using Actian are innovators in the financial services, healthcare, telecommunications, digital media and retail industries. The company is headquartered in Silicon Valley and has offices worldwide. www.actian.com



www.actian.com

2300 Geng Rd, Suite 150 Palo Alto, CA 94303
+1.888.446.4737 [Toll Free] | +1.650.587.5500 [Tel]



© 2016 Actian Corporation. Actian is a trademark of Actian Corporation and its subsidiaries. All other trademarks, trade names, service marks, and logos referenced herein belong to their respective companies. (WP08-0217)