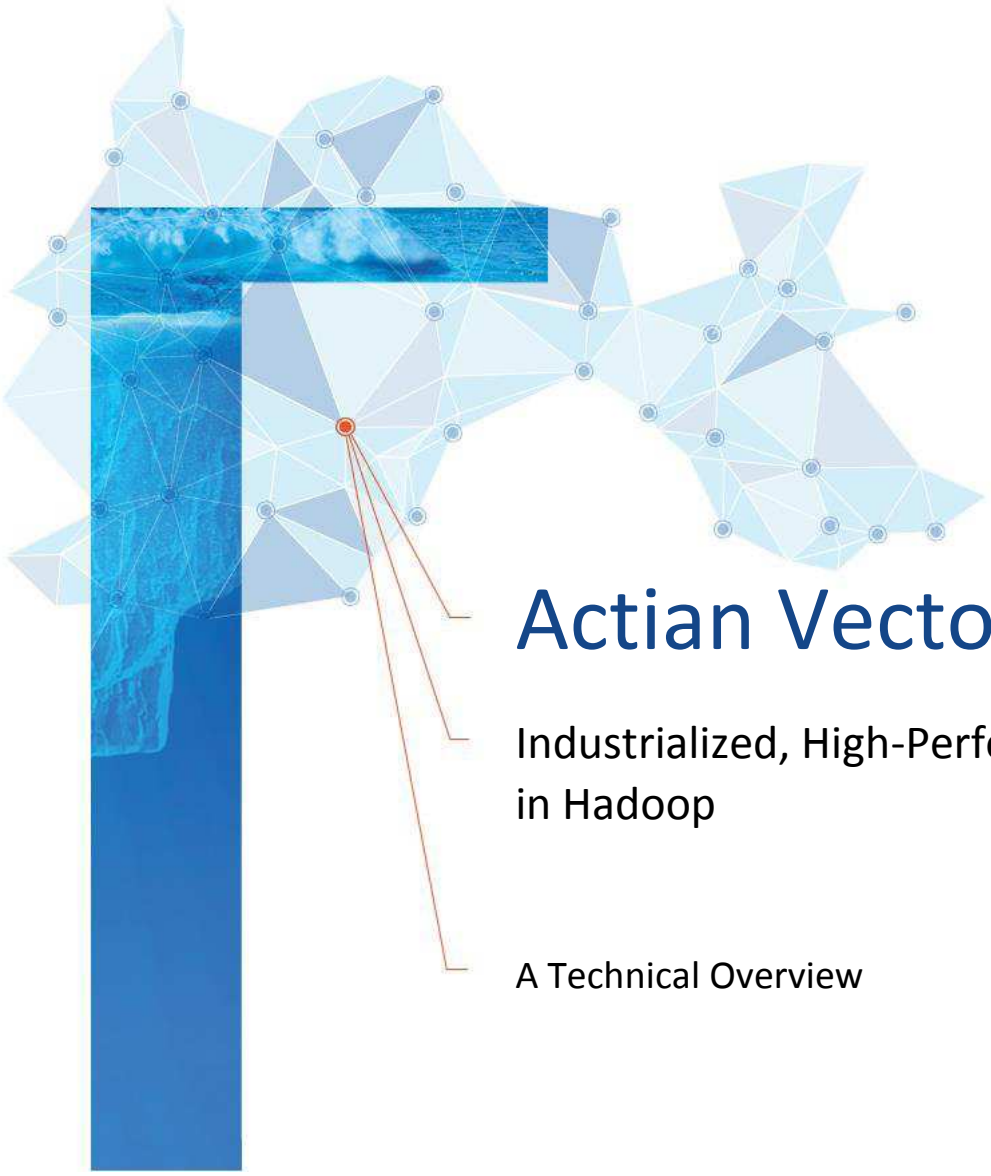

Actian Vector in Hadoop

Industrialized, High-Performance SQL in Hadoop

A Technical Overview

Contents

Introduction....3
Actian Vector in Hadoop - Uniquely Fast....5
Exploiting the CPU....5
Exploiting Single Instruction, Multiple Data (SIMD)....6
Utilizing CPU Cache as Execution Memory....6
Parallel Execution....6
Column-based Storage....7
Supporting Updates via Positional Delta Trees (PDTs)....7
Data Compression....7
Partitioned Table Support....8
Transparent DataNode Failover....8
Leveraging YARN for Resource Management....8
Actian Director for Ease of Administration....8

Actian Vector in Hadoop

Industrialized, High-Performance SQL in Hadoop

Hadoop has captured the imagination of those responsible for large scale data management and data analytics projects that are seen as game-changing to the enterprise. However, the need for rare and expensive skillsets, long and error-prone implementation cycles, lack of support for popular reporting and BI tools, and inadequate execution speed have led to a search for an alternative which combines all of the benefits of Hadoop with SQL – the world's most widely used data querying language. Actian Vector in Hadoop, also known as "Project Vortex", represents the next generation of innovation at Actian and employs a unique approach to bring all of the benefits of industrialized, high performance SQL together with Hadoop.

Introduction

Actian Vector is a proven relational database management system which has been extended to bring high performance data analytics to Hadoop. Actian Vector in Hadoop continues the legacy of cutting edge innovation to bring record beating performance, unique update capabilities, YARN resource management and distributed SQL query processing to Hadoop.

By bringing the raw performance of Actian Vector to the DataNodes in the Hadoop Cluster, Actian Vector is able to achieve performance levels never before seen with any SQL on Hadoop solution. Actian Vector 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 about the ability to query Hadoop data using standard BI and Reporting tools.

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

Hadoop Connector Approach: 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 and 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, their limited security features and lack of workload management. Solutions that fall into this category include Impala, Drill, HAWQ and IBM's Big SQL.

A new "SQL in Hadoop" category is required for Actian Vector in Hadoop. It 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. Actian Vector is a mature technology which has been hardened in the enterprise and includes support for localization and internationalization, advanced security features and workload management.

This paper explains how Actian Vector in Hadoop achieves extremely fast performance for typical data warehouse and data mart workloads in Hadoop. But don't just read this paper – experience Actian Vector in Hadoop in action in your own Hadoop environment. Get your copy of an evaluation version today. Contact us at eval@actian.com

Actian Vector in Hadoop - Uniquely Fast

To understand what's special about Actian Vector in Hadoop, let's review some of the key features that make Actian Vector unique.

Exploiting the CPU

Actian Vector takes advantage of powerful CPU features that most other databases can't or don't. During the past three decades CPU processing capacity has roughly followed Moore's Law. 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.

Actian 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. Actian Vector in Hadoop leverages these innovations and brings this unbridled processing power to the data nodes in a Hadoop Cluster.

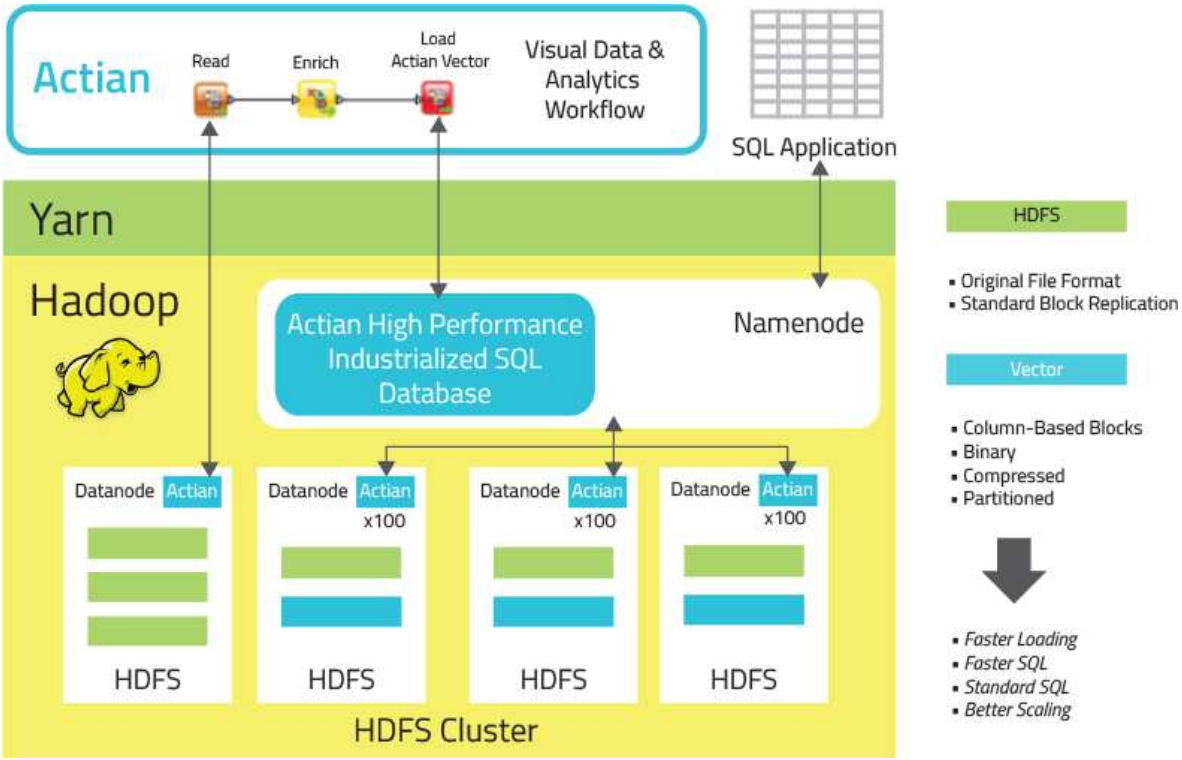

Exploiting Single Instruction, Multiple Data (SIMD)

SIMD enables a single operation to be applied on every entity in a set of data all at once. Actian Vector 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, Actian Vector 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 Actian 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 and slower over time. In addition, with more and 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, Actian Vector avoids the use of shared RAM as execution memory. Instead, Actian Vector uses the private CPU core and CPU caches on the Hadoop data nodes as execution memory, delivering significantly greater data processing throughput.

Parallel Execution

Actian Vector implements a flexible adaptive parallel execution algorithm and can be scaled-up or scaled-out to meet specific workload requirements. Actian Vector 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.

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

The NameNode serves as the single access point for client applications. It builds and distributes the parallel query execution plan and assembles the result sets before returning them to the client application.

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 accesses 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 a lot of unnecessary I/O for a Hadoop data warehouse workload. Actian's 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 Vector in Hadoop 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. Actian Vector in Hadoop addresses this challenge with 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 Vector in Hadoop 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 Vector 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 Vector 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 Vector 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 Vector 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. Actian Vector's decompression is far more efficient than alternative speed-optimized compression libraries such as LZOP that many other products have utilized.

Partitioned Table Support

Actian Vector in Hadoop 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 Vector in Hadoop can be attributed in part to elimination of all-to-all communications through the ability to execute joins locally on each DataNode.

Transparent DataNode 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 have them re-replicated to other (new) nodes.

Leveraging 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 version 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. This means that, once integrated with YARN, non-MapReduce based distributed applications, such as Actian Vector query workloads, can run as first-class citizens on Hadoop clusters, sharing resources side-by-side with MapReduce based applications.

The next release of Actian Vector will leverage YARN for dynamic resource management.

Actian Director for Ease of Administration

Actian Director provides a simple point and click interface for managing, administering and querying Actian Vector instances on premise and in the cloud. Available for both Windows and

Linux, Actian Director has been extended to include support for the new features included in Actian Vector in Hadoop including partitioned tables and HDFS browsing.

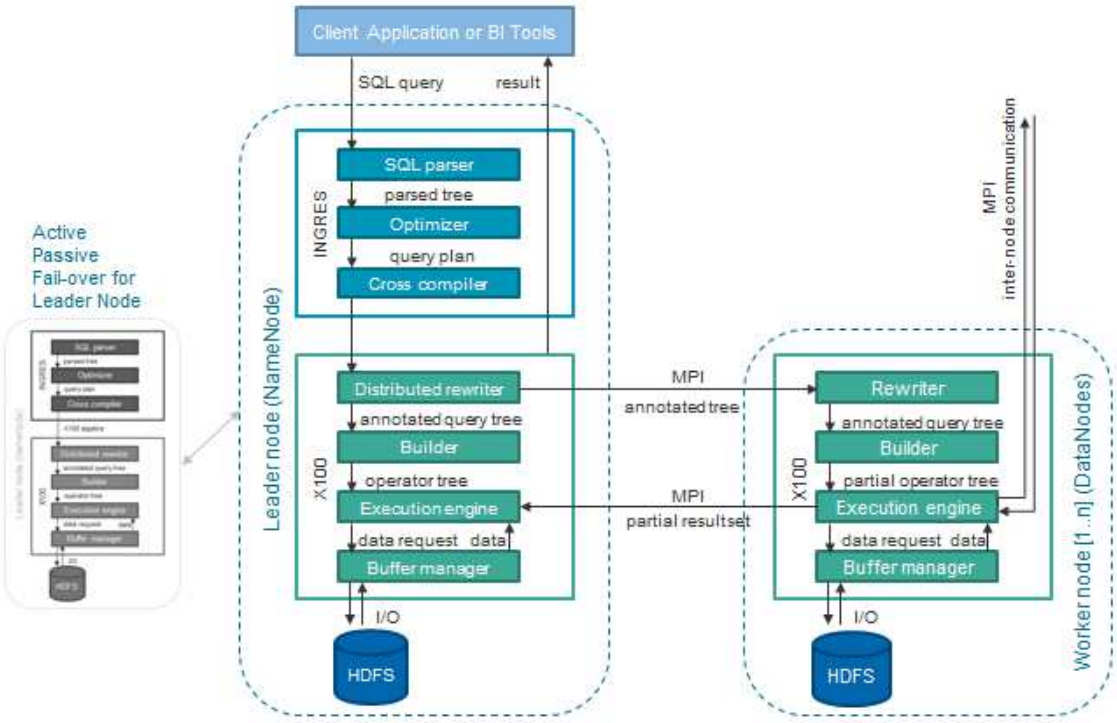

If you need to analyze large volumes of data on Hadoop and you don't want to take the risk of an expensive or lengthy implementation project, you should deploy Actian Vector in Hadoop. Implement an easy to deploy, easy to use, ANSI compliant solution and benefit from significantly better query performance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

You should try Actian Vector today but rest assured that there is more to come! Future versions of Actian Vector will not only introduce new functionality, leverage new Hadoop capabilities, but also continue to leverage CPU performance features and implement other optimizations to get absolute maximum query performance on Hadoop.

Use Actian Vector if you are looking for a relational database, supporting ANSI SQL and industry standard JDBC/ODBC interfaces that delivers extremely fast performance, is easy to use and is very cost-effective. Actian Vector delivers performance that is an order of magnitude faster than popular Hadoop databases.

About Actian: Accelerating Big Data 2.0

Actian transforms big data into business value for any organization – not just the privileged few. Actian provides transformational business value by delivering actionable insights into new sources of revenue, business opportunities, and ways of mitigating risk with high-performance in-database analytics complemented with extensive connectivity and data preparation. The 21st century software architecture of the Actian Analytics Platform delivers extreme performance on off-the-shelf hardware, overcoming key technical and economic barriers to broad adoption of big data. Actian also makes Hadoop enterprise-grade by providing high-performance data blending and enrichment, visual design and SQL analytics on Hadoop without the need for MapReduce skills. Among tens of thousands of organizations using Actian are innovators using analytics for competitive advantage in industries like financial services, telecommunications, digital media, healthcare and retail. The company is headquartered in Silicon Valley and has offices worldwide. www.actian.com

For more information or to request an evaluation version, contact us info@actian.com

www.actian.com

500 Arguello Street, Ste. 200, Redwood City, CA 94063
+1.888.446.4737 [Toll Free] | +1.650.587.5500 [Tel]