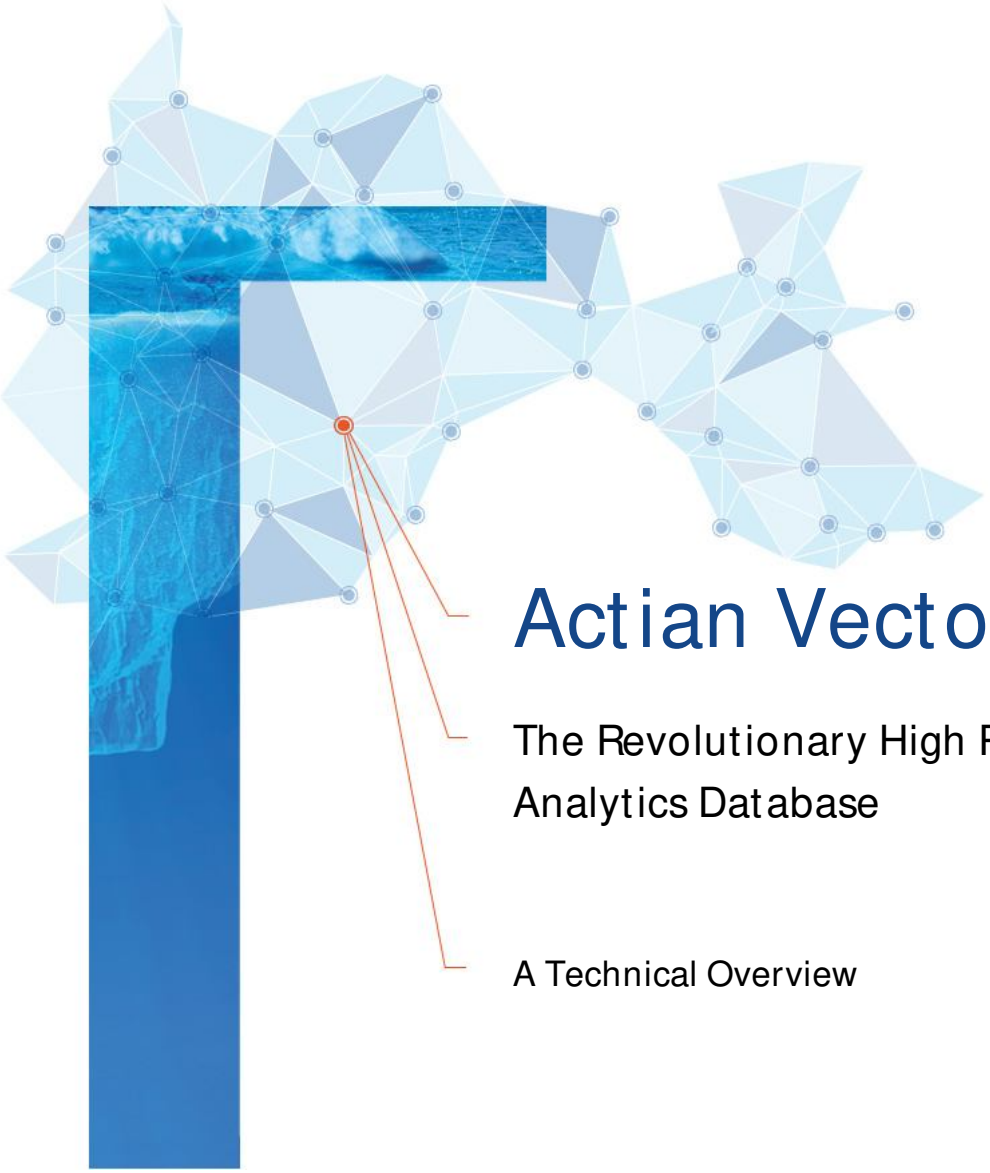

Actian Vector

The Revolutionary High Performance Analytics Database

A Technical Overview

Contents

Introduction

There is not a lack of data. Internet users and devices generate more and more each day. Companies and organizations recognize the need to analyze data in order to take action whether it is data generated by business processes or public data. Companies and organizations create data warehouses and data marts in relational databases to store and analyze terabytes of data – Big Data.

The market for relational database solutions for data warehousing or data marts has evolved rapidly over the last few years. Multiple purpose-built products are available for reporting, data analysis and Business Intelligence. Some product offerings are available only as a hardware and software combination – a data warehouse appliance – while others are software-only solutions that support a variety of hardware installations.

Actian Vector is relational database software for data analytics. Actian Vector exploits performance features in today's x86 CPUs that most other relational databases do not take advantage of. As a result, Actian Vector can process data much faster than most other relational databases. Much faster data processing performance opens up opportunities. Think not only about support for larger data sets, more users and more complex workloads, but also about the ability to directly query detail data when previously query performance would only be acceptable after extensive indexing and materialization of intermediate results. Faster performance significantly reduces the amount of lag-time until you can first look at results, and faster performance increases flexibility in the ways you can access your data.

But there is more. Actian Vector enables you to run a workload on a server when other databases require a much larger machine, a cluster of servers, or both, to achieve similar results. You can lower costs instantly by better utilizing your hardware, and also over time, since you don't have to carefully tune the system with hard-to-find experts.

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

Uniquely Fast – Exploiting the CPU

Actian Vector is unique because it takes advantage of powerful CPU features that most other databases don't. During the past three decades CPU processing capacity has roughly followed Moore's Law[1]. 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.

---

[1] Moore's law describes a long-term trend that the number of transistors that can be placed inexpensively on an integrated circuit doubles roughly every two years. See http://en.wikipedia.org/wiki/Moores_Law. Although Moore's Law specifically talks about the number of transistors it is casually used to describe technology improvements that double performance every two years.

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[2] 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.

Exploiting Single Instruction, Multiple Data (SIMD)

SIMD enables a single operation to be applied on a set of data 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. As a result, the CPU resources are used to perform the actual work.

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 as execution memory, delivering significantly greater data processing throughput.

Other CPU performance features

On an ongoing basis, the Actian Vector development team looks for ways to improve data processing performance using modern chip technology. For example, recent Intel®chips support hardware-accelerated string-based operations which are exploited by Actian Vector. Operations that benefit from the hardware-accelerated string based optimizations include selections on strings using wild card matching, aggregations on string-based values and joins or sorts using

---

[2] SIMD stands for Single Instruction, Multiple Data. Traditionally CPUs would process using a SISD model: Single Instruction, Single Data. For more information see http://en.wikipedia.org/wiki/SIMD.

string keys. However, not all modern CPUs support hardware-accelerated string-based operations and Actian Vector also works fine – just a little less optimally – if this performance feature is not available.

Leveraging Industry Best Practices

Various specialized data warehouse products use a number of well-known techniques to achieve fast performance. In general, because of the data-intensive nature of a data warehousing workload, most techniques focus on limiting and optimizing input/output (IO).

For Actian Vector – because of its dramatically higher per CPU core data processing power – to limit IO is an absolute requirement in order to achieve good data processing performance. Actian Vector implements industry best practices to limit IO while it introduces innovations to overcome some of the traditional weaknesses associated with these techniques.

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 is always retrieved row-by-row, even if a query only accesses a subset of the columns. This storage model works very well for On-Line Transaction Processing (OLTP) systems in which data is stored highly normalized, tables are relatively narrow, queries often retrieve very few rows and many small transactions can come through.

- Data warehouse databases are different:
- Tables are often (partially) denormalized resulting in many more columns per table, not all of which are accessed by most operations.
- 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 IO for a data warehouse workload. A 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.

Column-based databases have been available commercially for more than a decade. In addition to the benefit of data elimination when accessing fewer than all table columns in a query, an additional significant advantage of column-based storage is better data compression.

Hybrid Column Store

Actian Vector implements a hybrid column store. The term that is used in the research world for the type of storage Vector uses is PAX[3].

- By default, data is stored using a pure column-by-column approach.
- For tables that are indexed on more than one column, Actian Vector stores the indexed columns together in a single data block (but within the block, data is still stored column-by-column to optimize compression) assuming that indexed columns are typically accessed together.
- The user may choose to store data row-by-row if data allocation for column-by-column storage requires too much up-front data allocation. The choice for row-based storage can make sense for extremely wide tables or tables with relatively few rows.

Positional Delta Trees (PDTs)

Actian Vector implements a fully ACID[4] – 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 most column-based databases is incremental small inserts, updates or deletes (as opposed to large bulk data load operations). Actian Vector addresses this challenge with high-performance in-memory Positional Delta Trees (PDTs). Irrespective of the actual choice of data storage, Actian Vector uses PDTs to store small incremental changes (inserts that are not appends), as well as updates and deletes (except truncates).

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 disk. 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

Most relational databases support data compression and so does Actian Vector. It compresses data on a column-by-column, page-by-page basis using any one of the following algorithms or a combination of them:

- **Run Length Encoding** (RLE)[5]: a data value is stored as well as the number of subsequent values that are the same. This compression algorithm is very efficient on ordered data with relatively few unique values.

---

3 PAX stands for Partition Attributes Across. For more information visit http://www.pdl.cmu.edu/Database/index.shtml

4 ACID is an abbreviation that stands for Atomicity, Consistency, Isolation, Durability–a set of properties that guarantee database transactions are processed reliably. For more information visit http://en.wikipedia.org/wiki/ACID.

5 See http://en.wikipedia.org/wiki/Run-length_encoding.

- **Patched Frame Of Reference** (PFOR): a base value is determined per data block and other values in the same block are encoded by storing the difference with the stored value using as few bits as possible. This is beneficial as the range of the actual data is typically much smaller than the range of a used data type. What makes PFOR special compared to similar solutions found in other products is the treatment of outliers. For example, if 99% of values are in range 0–255, and 1% of value is very large (e.g. around a million), then with PFOR the vast majority of the data will be stored using only one byte, while other solutions would use 2.5 bytes.
- **Delta encoding on top of PFOR**: in order to reduce the values of the integers with PFOR, it is sometimes more efficient to store the delta from the previous value. This can be very efficient on ordered data.
- **Dictionary encoding**: stores pointers to a dictionary of unique values. This algorithm is very efficient for a limited number of very frequently occurring values.
- **LZ4**: detects and encodes common fragments of different string values. It is particularly efficient for medium and long strings.

The algorithms Actian Vector 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.

Actian Vector's innovative use of data compression

In order to improve IO 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 disk. 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.

Actian Vector automatically chooses the most optimal compression on a page by page basis and pages are large. I.e. per column – multiple pages – 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.

Storage indexes

Actian Vector automatically maintains a storage index per column storing minimum and maximum values for the data block. The storage index is very efficient in determining whether a database block is a candidate block for a particular query either because of explicit filter criteria or implicitly as a result of processing table joins.

In extreme cases, the storage index provides the same benefit as data partitioning does for other databases without the overhead of multiple database objects or having to design and maintain a partitioning strategy.

Parallel execution

Almost all relational databases support some means for a single operation to take advantage of multiple CPU core resources. For some databases, particularly the pure Massively Parallel Processing (MPP) databases, the use of multiple CPU cores is a mandate and virtually every operation uses all CPU cores in the system. Other databases use some form of a shared architecture and therefore support a wider range of possible degrees of parallelism.

Actian Vector implements a flexible adaptive parallel execution algorithm. Actian Vector can execute statements in parallel using any number of CPU cores up to the number of cores in the server, but if many operations run concurrently then parallelism is automatically reduced to make optimum use of the available system resources without overloading the server.

Actian Vector Use Cases

Actian Vector provides relational database software that takes analytic data processing to a new level. With it, you can now achieve amazing performance with a simple, ANSI compliant relational database – something that was previously only achievable with proprietary OLAP databases and/or following lots and lots of careful design and tuning using complex features.

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 faster than popular in-memory type databases without even having to load all data in memory and without the hard limit of available memory. The diagram below shows a number of areas where you should consider Actian Vector.

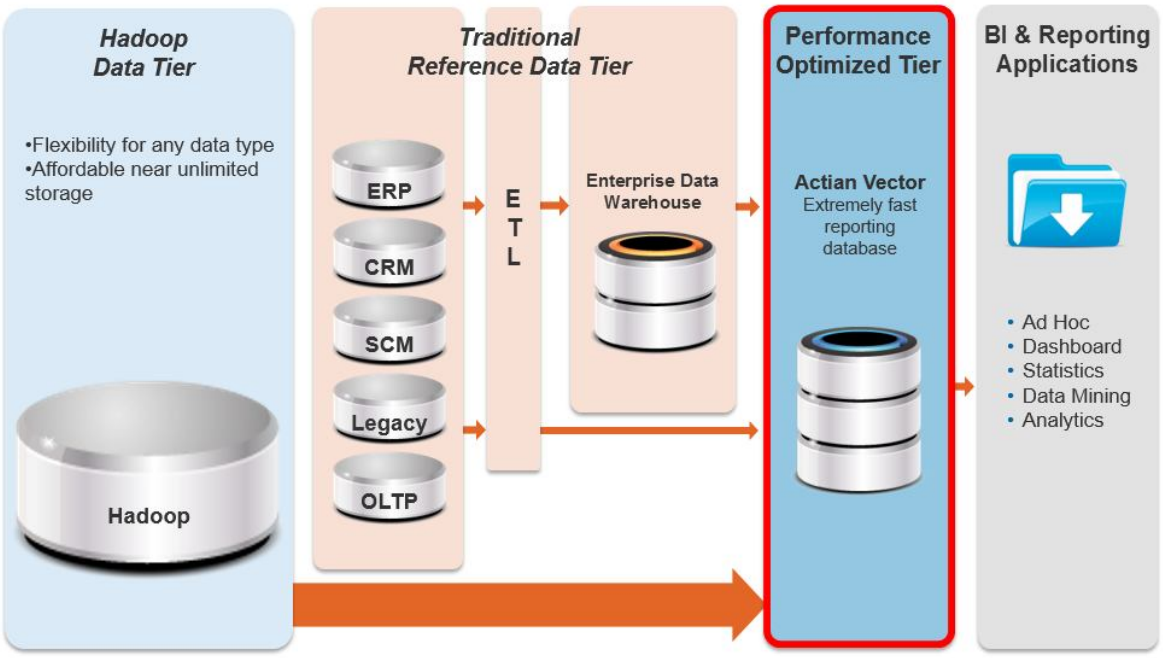

Figure 1. Cooperative Analytics with Actian Vector

Following are examples of data-intensive Actian Vector use cases.

Financial services

A number of organizations in financial services chose Actian Vector. The Rohatyn Group, a Wall Street-based hedge fund focusing on emerging technologies, replaced a home-grown, in-memory database with Actian Vector in order to continue to deliver at least as good in-memory-like performance while not being limited to the total amount of memory. The analysts using the system had expressed a desire to query historical data as well as current positions and the data volume was simply too large to store cost-effectively in memory.

Actian Vector provides the in-memory performance, but now hundreds of millions of rows containing historical data are stored on-disk.

Type equation here.*"For the past 20 years, I've been searching for the killer database that would fulfill most of our intense data processing needs and with the discovery of Actian Vector, that search is now over - this database is in a class of its own. Right out of the box, Actian Vector lets us effortlessly plow through millions and millions of rows of data with infinite width and depth and without the need for new expensive hardware, complicated schemas, explicit indexing, pre-aggregation, or specifically hand-crafted DBA-tuned SQL. The Actian leaders and technologists have performed a miracle here."*

— **Warren Master,** CTO, The Rohatyn Group

Retail

Sheetz is a $5 billion convenience store business with a reputation for progressive marketing and fierce competitiveness in the marketplace. From day one, company executives recognized the value of having a finger on the pulse of what consumers want from a convenience store. As the business grew, this became more of a challenge. By deploying Actian Vector, Sheetz gained the ability to analyze a far more comprehensive set of data (more than three billion rows), returning query results in seconds. It offered performance improvements of as much as 70X over conventional technology by utilizing the latent processing power in the company's existing hardware infrastructure, with the added benefit of reduced operational costs. And, Actian Vector enabled Sheetz to double its access to historical data and be ready for the expected growth over the next few years.

"Data growth is occurring at record rates. Based on our experience so far, we are impressed with the results of Actian Vector..."

— **Jarrid Magalich**, Technical Services Manager at Sheetz

Social media

Many social media websites are extremely popular and generate vast amounts of data about their users. NK (http://nk.pl) is a large social media site in Poland, twice as large as Facebook in Poland.

Social media companies often use advertising to monetize user behavior. However, user behavior changes and changes in behavior warrant action. Prior to using Actian Vector, the Product Managers at NK would have to wait days or weeks for their queries against the vast amounts of data to complete. They were answering business requests for data with workaround solutions built on MySQL databases with huge queries. NK implemented a new solution that collects data from various sources, including Hadoop, and imports it into Actian Vector for its ultrafast processing performance. Actian Vector has enabled faster report generation, improved user experience via dynamic dashboards, simplified queries, and improved access to harmonized data.

"We looked to solutions from other vendors with analytic databases, but selected Actian Vector for its superior performance and cost-effective model."

— **Edward Mezyk**, Senior Project Coordinator in NK Research and Data-warehouse Division

Digital Media

Edo interactive, a B2B electronic marketing firm, provides 120 million offers a month and more than 25 million transactions a day - producing as much as 50 terabytes of data. They needed low latency interactive SQL analytics in support of various business user groups. By deploying Actian Vector, they were able to analyze and visualize patterns in sub-1-minute queries over terabytes of highly structured data. Now, business users have self-service analytics on terabytes of data.

"With Actian Vector, edo Interactive team gained rapid access to data along with support for analytic queries that we were unable to experience before."

—**Tim Garnto,** vice president of product engineering at edo Interactive

Transportation and Distribution

Timocom is the leader in European freight exchange. On a daily basis Timocom brings some 85,000 users and 300,000 international cargo space and freight offers together on its web portal. One of the challenges for freight carriers is fraud and theft. Timocom selected Actian Vector to help monitor for criminal activity and analyze user behavior.

"Our database in its existing form had reached its limits. Generally, we have already seen up to hundredfold of the inquiry speed in our initial tests without having done any optimization of the tables or inquiries."

— **Ingo Klose**, Manager Business Intelligence at TimoCom

Clinical Research

CTSU (the MRC/Cancer Research UK/BHF Clinical Trial Service Unit & Epidemiological Studies Unit of Oxford University) primarily studies the causes and treatment of chronic diseases such as cancer, heart attack or stroke (which, collectively, account for most adult deaths worldwide). Vast data volumes are analyzed to look for a needle in a haystack. CTSU selected Actian Vector to perform analyses of these massive data sets in minutes rather than hours or weeks.

"Without Actian Vector, we simply would not be able to process this information, without having to wait days or weeks for each output."

— **Alan Young**, Director of Information Science, CTSU

Conclusion

Actian Vector is the first relational database that focuses on processing efficiency in modern CPUs. Its vector-based processing as well as other optimizations directly take advantage of improvements in modern chips. Actian Vector is available on cost-effective x86-64 Linux and Windows platforms.

In order to maximize performance, the entire underlying database architecture is designed to eliminate any potential bottlenecks that would limit CPU processing. Column-based storage, data compression and smart storage indexes are all means to achieving this goal. In addition, parallel execution can squeeze the absolute maximum performance out of a system.

If you need to analyze large volumes of data and you don't want to take the risk of an expensive or lengthy implementation project, you should deploy Actian Vector. Implement an easy to deploy, easy to use solution and benefit from significantly better query performance than other relational databases. Actian Vector is the foundation for revolutionary performance gains in database processing. 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, but also continue to leverage CPU performance features and implement other optimizations to get absolute maximum query performance. Actian Vector enables users to gain timely insights into big data and take action.

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.

For more information and to request an evaluation version, go to www.actian.com