



Squeezing More Value Out of Data Warehouses through the Power of Database Compression

Using the power of data compression to lower data warehouse ownership costs and accelerate performance.

Copyright Vertica Systems Inc. March 2010

Table of Contents

Table of Contents	2
Executive Summary	3
Introduction	3
Why is Compression important?	4
How does Compression work in Vertica?	4
Column Oriented Storage	4
Sorting	5
LZ Compression	6
Encoding	6
How much compression can I expect for 'my' data?	7
Customer Success Stories	8
Benefits of Vertica Compression	8
Query Performance	8
Saving on Storage Costs	9
Saving on Operational Costs	9
Security	9
Re-purpose storage saved by compression	9
Additional Resources	10
About Vertica Systems	10

Executive Summary

In a world of growing data volumes and shrinking IT budgets, it is critical to think differently about the efficiency of your database and storage infrastructure. The Vertica Analytic Database is a high-performance, scalable and cost-effective solution that can bring dramatic savings in hardware, storage and operational costs. With its columnar storage architecture and compression-aware query engine, Vertica can speed up queries to run orders of magnitude faster (50x-300x based on customer benchmarks), and at the same time shrink database storage by 50-90%. Watch as more and more users within your enterprise are empowered by high performance data analytics, and without breaking the bank!

“With tens of terabytes of raw SNMP data captured from Comcast’s cable network for analysis, compression is a big factor for our network performance monitoring data warehouse. Disk I/O is always the performance bottleneck, and writing 90% less data to disk means less IO and more cost savings on the storage side.”

Brian Harvell
Sr. Director Network Operations
Comcast

Introduction

The amount of business data is exploding at a dramatic rate. No matter the industry or size of the business, massive amounts of data are generated every day. Data analytics are being applied to much more than traditional transactional data – new classes of applications are emerging around analysis of stock tick data, call-detail records, web click stream data and so on. Another trend that is emerging is that data is being retained for longer periods of time – either because more history provides more accurate analysis or because compliance laws force retention.

In the current economic climate, it is not a stretch to think that the enterprises that are savvy about the way they store, analyze and retain their data will be the ones to survive and thrive. For example, better data analysis around securitized mortgage assets could help financial companies more accurately quantify their exposure to risky sub-prime assets. The quality of this analysis depends entirely on how much historical data is available to validate the risk model and how efficiently the new data can be analyzed against it. This will no doubt lead to increased storage needs and greater pressure on the existing database infrastructure to perform. As data storage needs increase, so does the cost of storage hardware and the cost of managing the storage.

The Vertica Analytic Database presents a high-performance, scalable and cost-effective solution to these challenges. With its innovative compression technology, Vertica greatly reduces the cost-per-terabyte of data storage. At the same time, Vertica’s compression-aware query engine leverages the advances in multi-core CPU technology to fundamentally switch the bottlenecks in data analytics from the disk I/O tier to the compute tier. Regardless of whether you have a large enterprise data warehouse performance problem, a subject-area data analysis problem or increasing data retention requirements, Vertica can reduce the total cost of ownership (TCO) of your analytic databases and data warehouses while giving more users within your enterprise the ability to analyze more data faster.

Sounds like a free lunch? Read on to find out why compression is so important today, the unique benefits of Vertica’s compression technology, how it is designed to be on the “right-side of Moore’s law” and why a large number of companies are including Vertica as a strategic component within their data warehousing infrastructure.

Why is Compression important?

Compression by definition reduces storage space occupied by the database. So what's new? Hardware trends favor doing more with compression. With advances in processor technology, CPUs are getting faster and cheaper all the time (according to Moore's Law!). On the other hand, disks are only getting bigger in capacity but not much faster. Hence, a database architecture like Vertica's that reduces disk bandwidth requirements by smart use of compression, potentially at the cost of increasing CPU usage is on the right side of Moore's Law – as CPUs get better and cheaper, systems running Vertica will automatically benefit from these advances. For an excellent discussion on disk versus CPU trends, the reader is referred to a recent blog post by well-known database authority, Dr. Dave Dewitt of Microsoft. (<http://www.databasecolumn.com/2007/10/cpu-trends-like-disk-trends.html>)

In the next section we look under-the-hood to understand the technical reasons behind Vertica's superior compression technology.

How does Compression work in Vertica?

The Vertica Analytic Database uses a fundamentally different column-oriented storage architecture that lends itself to significantly enhanced compression. Further, the Vertica query engine is designed from the ground-up to take advantage of the compressed storage format and the emerging multi-core CPU architectures. This smart use of compression can result in out-of-the-box speed-ups of over a magnitude on query performance, as well as significant storage footprint reduction over the raw data size. Let us look at each of these factors in detail.

Column Oriented Storage

Unlike traditional databases that store data in a row-oriented fashion, Vertica stores data in a column-oriented manner, as illustrated in Figure 1. By storing each column separately, Vertica can speed-up queries by retrieving from disk only those columns requested by the query. In a typical call detail record (CDR) or stock tickstore application with 300 columns, most queries retrieve only 8 or 10 columns. Thus, column-orientation alone can make a significant dent in the I/O bandwidth requirements of the database. Unlike Vertica, row-oriented databases pay a heavy price for having wide tables, requiring auxiliary structures like materialized views (MVs) and bitmap indexes to speed up query performance. The added cost of storage is tremendous – often times, storage occupied by MVs can be 2 to 3 times that of the base data. It is well-known that these tuning structures cannot be efficiently maintained if data is loaded continuously rather than once a day as a batch job. Further yet, if a surprise query is thrown in for which no MV was designed, performance degradation for that query can be dramatic. There are no such trade-offs or surprise killer queries in a column-oriented Vertica database – your baseline performance is guaranteed to be better for the reasons outlined above.

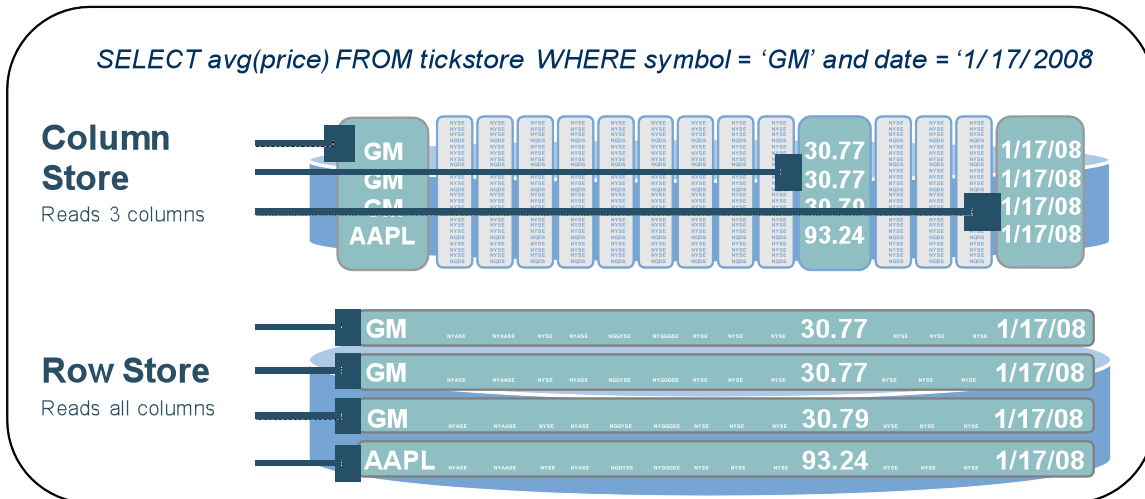


Figure: 1 Comparing Vertica's column-oriented storage to traditional row-oriented storage

Sorting

Fine – but what do columns have to do with compression? Well, it turns out that column storage compresses better. Why? The data values within a single column have identical data types and there is likely to be much more repetition within a column than across multiple columns. Compression techniques generally work best when there is some similarity within the data and so when data is stored in columns, a much higher compression ratio can be achieved. All row-store databases, including modern MPP variants of open source databases (such as clustered MySQL or Postgres) suffer from this inherent disadvantage with respect to compressibility. Also, because traditional row-oriented databases were designed mainly for handling transaction processing (OLTP) workloads and thereby optimized for updates rather than queries, disk blocks must have free space left for updates, which further limits the effectiveness of compression. Finally, even if reasonable compression is achieved on initial load, it degrades as more data is added to the database – the problem is acute if the database is continuously loaded, which is increasingly becoming a critical requirement for many analytic applications.

Vertica solves the issue of maintaining compression through loads in an interesting way. Rather than storing data in the order in which it is loaded, Vertica stores data in an order which is closely aligned to expected query patterns. This is done using an internal storage structure called a projection, which is simply a subset of columns from the table, stored in a particular sort order. Each projection consists of a number of containers of sorted data stored in a densely packed manner i.e. there is no free space. As more data gets added, Vertica transparently reorganizes storage to achieve optimal number of compressed containers for fastest query performance. This storage structure is perfectly aligned to the nature of the analytic workload, which consists mainly of continuous loads and queries. For example, in a Tickstore schema where most queries ask questions regarding a stock symbol or over a period of time, why not store the data by symbol and by date? The sorted data results in “like” data being clustered together on disk, resulting in even better compression. Note that, these sort orders are not auxiliary structures but merely Vertica's internal storage organization – in Vertica you do not pay the price of storing the base tables – you only store the data the way most useful for queries.

LZ Compression

Traditional compression techniques like Lempel-Ziv (LZ) are opaque, which means it is not possible to look inside a compressed chunk of data and find the original data without decompressing it first. Most relational databases that claim to do compression use some variant of LZ compression (such as LZO, LZW, etc). As discussed earlier, row-oriented storage is inherently inefficient for compression. Also, in these databases, compression is an after-thought rather than fundamental to the engine design. So the best they can do is to encapsulate it in the storage layer of the database – as soon as it is read from disk, it is uncompressed for the rest of the engine to consume. The result is performance improvements only up to 30% and similarly modest space savings, much of which may be lost by adding auxiliary data structures like materialized views. Vertica uses LZ compression as well, but only as a last resort for random floating point data. The technique primarily responsible for superior compression in Vertica is **encoding**, which is discussed in the next section.

Encoding

Compression techniques such as Run-Length-Encoding (RLE), Delta Value, Block Dictionary etc. have the property that the original data can be retrieved without un-encoding the data. These “data aware” compression techniques, also known as Encoding can provide significantly superior compression compared to LZ variants. The Vertica query engine and all database operations have been designed to take full advantage of a variety of encodings. Operations such as joins, aggregations, sorting, data transfers across the network - all considered very expensive operations in traditional architectures, are significantly faster in Vertica because less data is involved. In addition to the well-known encoding techniques, Vertica uses optimized integer and floating point compression algorithms as well as compound encoding algorithms that reduce storage footprint and enhance overall system performance. Vertica has over a dozen compression schemes that can be employed depending on the data. A surprisingly large class of data is amenable to encoding - even seemingly random data like stock prices have a fairly narrow range of variation within say, a trading day, which makes it very well suited to Delta Value encoding. Some encoding schemes also have the curious property that the compressed size may not grow at all or grow sub-linearly with addition of more data – for example, it costs the same amount of storage (roughly 1KB!) to store a sorted and RLE column with data for the 50 US states, regardless of whether the table is 1000 rows or 100 million rows.

Another great benefit of encoding is in storage of NULL values which can lead to a huge amount of wasted space in row-oriented architectures. Because a NULL value can be treated like any other data value, depending on the data type of the column, different encoding schemes may be used to represent it. SQL operations that have special semantics for NULLs, can be efficiently handled while preserving the data in the encoded form.

The right compression technique to use is very much dependent on the data. Rather than use a one-size-fits-all compression scheme for all kinds of data, a Vertica implementation will be smart about how the data is laid out, sorted and compressed, giving you the best possible performance for every dollar spent on storage. The next two sections will present real-world compression numbers seen by our customers.

How much compression can I expect for 'my' data?

While a number of database vendors claim anywhere from 30:1 to 2:1 compression, in reality, compression is very dependent on data and how it is loaded. Analyzing applications across many industries we have found databases are compressible to different extents.

In financial services, stock market trade data for all US exchanges includes 250 days of data and 10,000 instruments each year. At 100 million trades per day (a conservative number) a dataset that records the date and time, instrument, price and volume will use between 3GB and 10GB (uncompressed) per day depending on the representation (e.g. binary vs. ASCII). By one year, this raw data will be anywhere from 1 to 3TB and using an off-the-shelf LZ algorithm may compress it by 2:1.

In the telecommunications space, call detail records include information describing the call path, times, features, and switches. An average record may be 550 bytes long and a regional provider may record 500 million per day resulting in uncompressed source data of 275GB per day or 100TB a year (calls are made 365 days a year). This data, also using off the shelf LZ, compresses by 5:1. These compression factors relative to raw data can vary dramatically depending on whether the data is preprocessed, for example, encoding long names into ids, changing timestamps with time-zones into GMT or replacing empty strings with nulls.

The Vertica customer base spans many industries and applications, so we have been able to assess compression with a large variety of real world datasets. We see the following as typical compression ratios, relative to raw ASCII delimited data (e.g. comma separated values):

- Call detail records (CDR) – 8:1 (87%)¹
- Consumer data 30:1 (96%)
- Marketing analytics 20:1 (95%)
- Network logging 60:1 (98%)
- Switch-level SNMP 13:1 (92%)
- Trade and quote exchange – 5:1 (80%)
- Trade execution auditing Trails – 10:1 (90%)
- Weblog and click-stream 10:1 (90%)

Your mileage may vary depending on the nature of your data set, however Vertica welcomes you to evaluate the product on your data set and determine the compression savings for yourself.

¹ Note: A ratio of 2:1 is equivalent to 50% compression while 30:1 is equivalent to 96.6% compression. For a ratio of R, the percent compression is $1 - 1/R$. Correspondingly the ratio for a given percentage P is $1/(1 - P)$.

Customer Success Stories

The power of compression in Vertica is best evidenced by its effectiveness for our customers. Below is a sampling of some specific customer use-cases.

Comcast uses Vertica to capture and analyze data generated by millions of components in their cable communications network for performance monitoring purposes. Vertica runs on a 5-node cluster of HP ProLiant DL380 servers, each with 1TB of internal storage (for a total hardware cost of approximately \$50,000 USD). With a measured compression rate of 11:1, Comcast expects the 5-node system to scale past 20TB of raw data (including built-in 2x data redundancy for fault-tolerance). They can simply add new nodes to the grid to extend that in the future. The system currently loads 15GB of new data per day on a continuous basis and answers queries in less than 1 second.

A large US telecommunications provider runs Vertica on a cluster of 6 HP ProLiant DL380 servers with 1TB internal storage per node (total hardware cost is under \$60,000 USD) for a call detail record (CDR) data warehouse. With an initial raw data volume of 5TB and a continuous trickle load of 30MB of new CDRs per minute and a compression ratio of 8:1, the current hardware can sustain database growth up to 20TB of user data including 2x redundancy. With tens of users accessing the system simultaneously, 80% of the queries take less than a second.

A global financial investment firm uses Vertica for a trade analysis application on a 10-node cluster using Dell servers with 2TB internal storage per node. Daily load rates are 200GB. Compression ratios observed are 6:1, and the system is expected to store up to 40-50TB of data, including data redundancy for fault tolerance.

These are just a few examples of Vertica customers who have been able to achieve a very low cost per terabyte of storage by effective use of compression. The examples also show that compression does not prevent high loading speeds from being achieved. In fact, In December 2008, ETL software vendor, Syncsort, HP and Vertica teamed up to break the world record for data warehouse data loading. It took 57 minutes 21 seconds to load 5.4TB of TPC-H data (generated retail sales data) into Vertica by the Syncsort ETL software running on an HP BladeSystem C7000, which beat the previous record set by Microsoft (2.36TB in 1 hour). When the load was complete, the size of the Vertica database holding the 5.4TB of data was just 1.54TB (a compression ratio of 3.5:1).

Benefits of Vertica Compression

Compression in Vertica directly contributes to reducing the TCO of a Vertica system. Let's review some of the benefits of Vertica compression.

Query Performance

Query performance is improved significantly on account of reduced disk I/O and the ability of the engine to operate on encoded data. This means you can achieve the great performance on much less hardware, directly contributing to a lower TCO.

Saving on Storage Costs

The reduction in storage footprint directly leads to a lower storage cost per terabyte. It is instructive to note that the typical size of a Vertica database is much smaller than the size of raw data, unlike other databases which experience bloat due to auxiliary access structures.

Saving on Operational Costs

Compression in Vertica can have other operational benefits. Storage management is more efficient – for example, fewer storage components, less storage performance tuning and less time spent on backups. Capacity planning is simplified - unlike other databases, you do not need to worry about accounting for storage of auxiliary data structures that may be required in the future for performance tuning. To determine how much storage you will need, you simply do some experiments on a sample of your data and use that as guideline.

Security

The columnar and compressed representation in Vertica makes it extremely difficult to reverse engineer the Vertica data format on disk. Column values across multiple rows get mixed together into an encoded representation. Often more than one encoding and compression scheme may be combined for the same column. Different columns may be compressed differently. As a result, compression in Vertica in effect provides obfuscation of data as well. It can be considered as a simple but highly effective form of encryption that does not negatively affect performance.

Re-purpose storage saved by compression

The storage reclaimed by compression can be gainfully employed by storing more data in the database (e.g., more historic data) or by having Vertica automatically replicate data (by creating many different overlapping sub-sets of columns with different sort orders (*projections*)) to provide higher redundancy. When the entire system is up, the multiple projections are used to service different classes of queries. When one of the nodes is down, the remaining projections can be used to keep the system running with minimal performance impact. The storage saved by the use of compression can be employed to create additional projections - which could potentially also service more classes of queries.

With all these benefits, it is clear that no matter what your immediate pain point is, Vertica can provide you with a strategic advantage in today's tough economic conditions by enhancing query performance, while at the same time reducing your IT infrastructure and operational costs.

Additional Resources

The Vertica Analytic Database supports SQL and integrates with 3rd-party ETL, analytic and BI reporting tools and applications via JDBC, ODBC and specific language bindings. Therefore, using all your existing SQL knowledge and technology, a Vertica database can be very quickly created and loaded with data.

If you would like to learn more about the Vertica Analytic Database or experience the power of compression for yourself, please visit www.vertica.com and request an evaluation copy of Vertica. Below is a list of links about Vertica technology you may find interesting:

Gartner on Vertica and EDWs	www.vertica.com/gartner	Watch a recording of Don Feinberg of Gartner explain why supplementing an EDW with column-DB based data marts improves the ROI on EDWs
Vertica Resource Library	www.vertica.com/resourcelibrary	White papers, demos, webcasts, system requirements
Vertica Benchmarks	www.vertica.com/benchmarks	See customer-submitted cost and performance comparisons between Vertica and other databases
Vertica for the Cloud	www.vertica.com/cloud	Get a Vertica database instance provisioned instantly on the Amazon Cloud and use it on a month-to-month basis
Vertica Customers	www.vertica.com/customers	See who's using Vertica
Request a Vertica Evaluation	www.vertica.com/download	Request a free evaluation copy of the Vertica Analytic Database to download and install

About Vertica Systems

Vertica Systems is the market innovator for high-performance analytic database management systems that run on industry-standard hardware. Co-founded by database pioneer Dr. Michael Stonebraker, Vertica has developed grid-based, column-oriented analytic database technology that lets companies of any size store and query very large databases orders of magnitude faster and more affordably than other solutions. The Vertica Analytic Database's unmatched speed, scalability, flexibility and ease of use helps customers like JP Morgan Chase, Verizon, Mozilla, Comcast, Level 3 Communications and Vonage capitalize on business opportunities in real time. For more information, visit the company's Web site at <http://www.vertica.com>.